Energy Efficient Performance: The New Frontier

Dileep Bhandarkar
Architect At Large
Digital Enterprise Group
Intel Corporation

October 10, 2006

Legal Disclaimer

- INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT.
- Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.
- Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.
- Intel may make changes to specifications, product descriptions, and plans at any time, without notice.
- The processors, the chipsets, or the ICH components may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available upon request.
- All dates provided are subject to change without notice. All dates specified are target dates, are provided for planning
purposes only and are subject to change.

- Intel and are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- * Other names and brands may be claimed as the property of others.
-

1993

2003

The GHz Era

intel

Increasing Performance

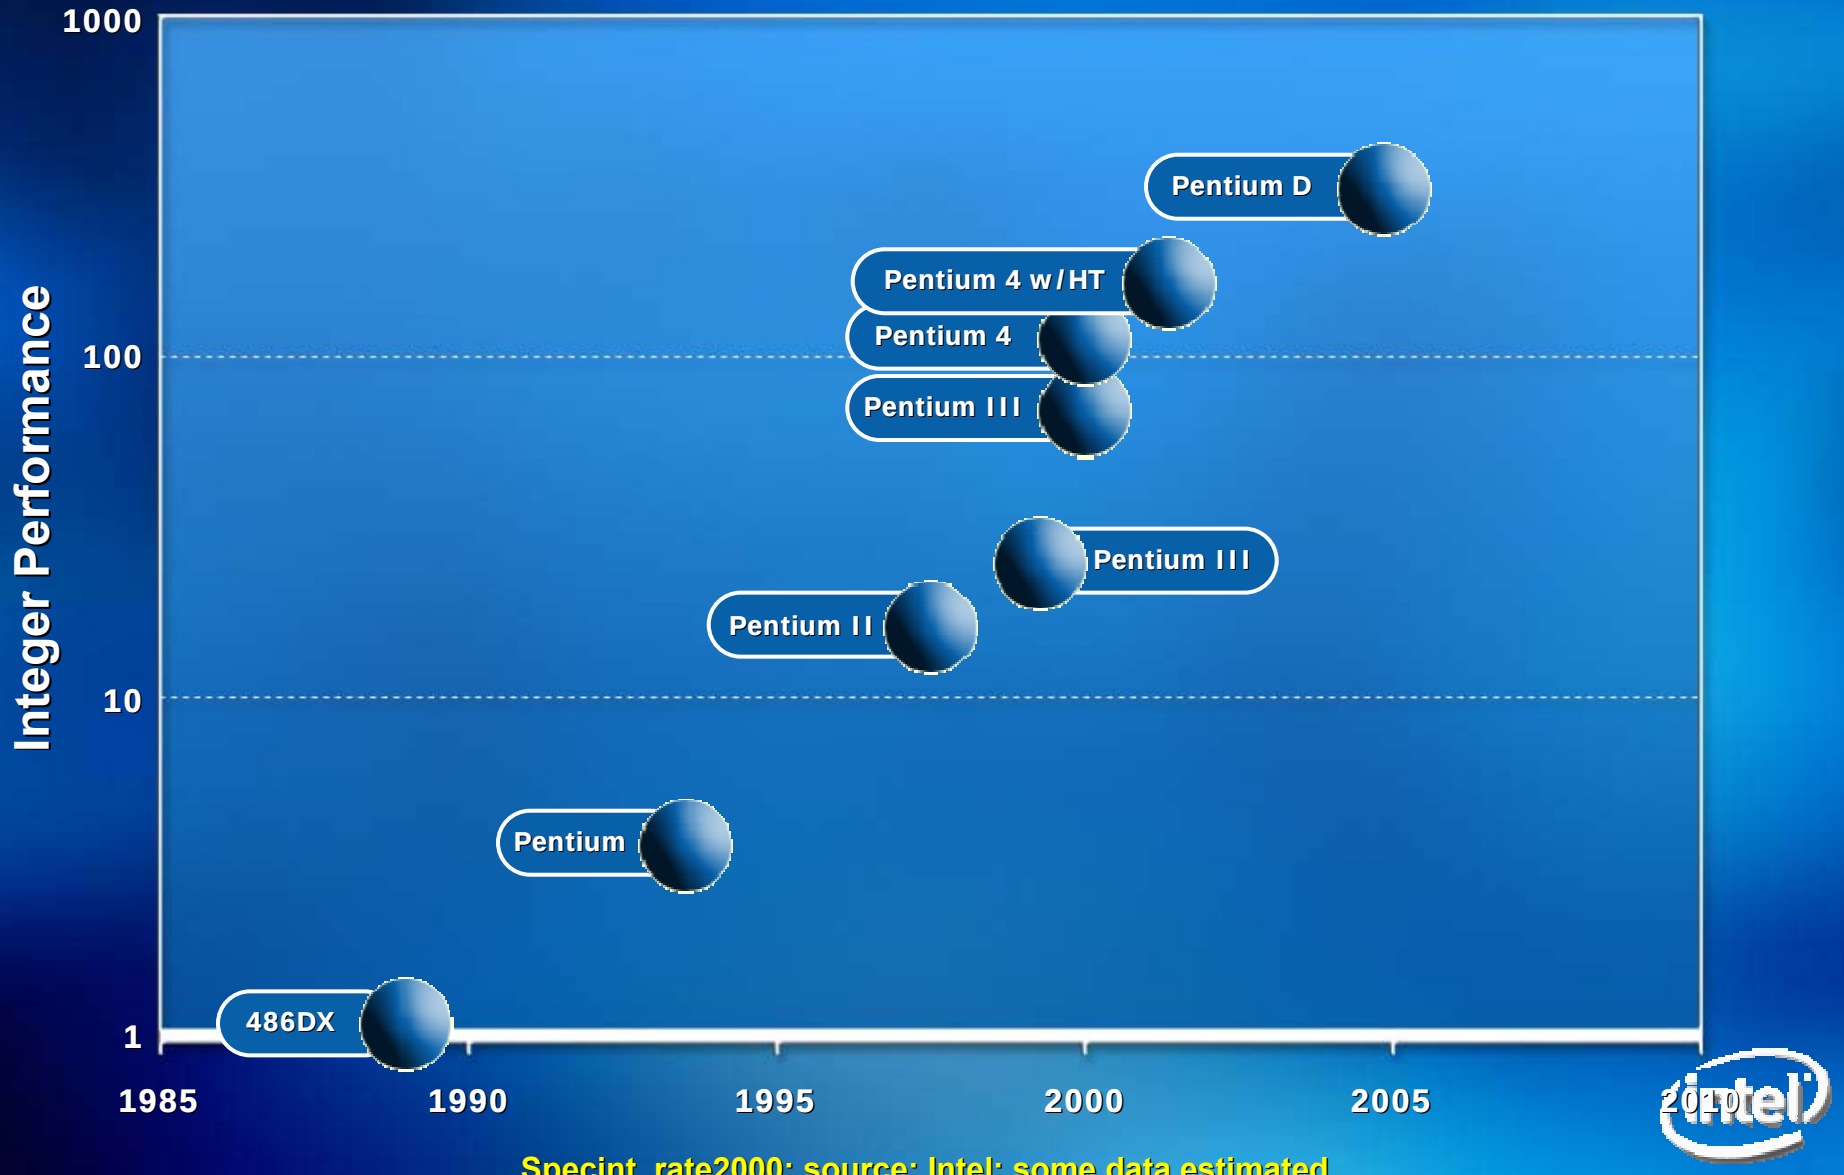

Specint_rate2000; source: Intel; some data estimated.

Increasing Power

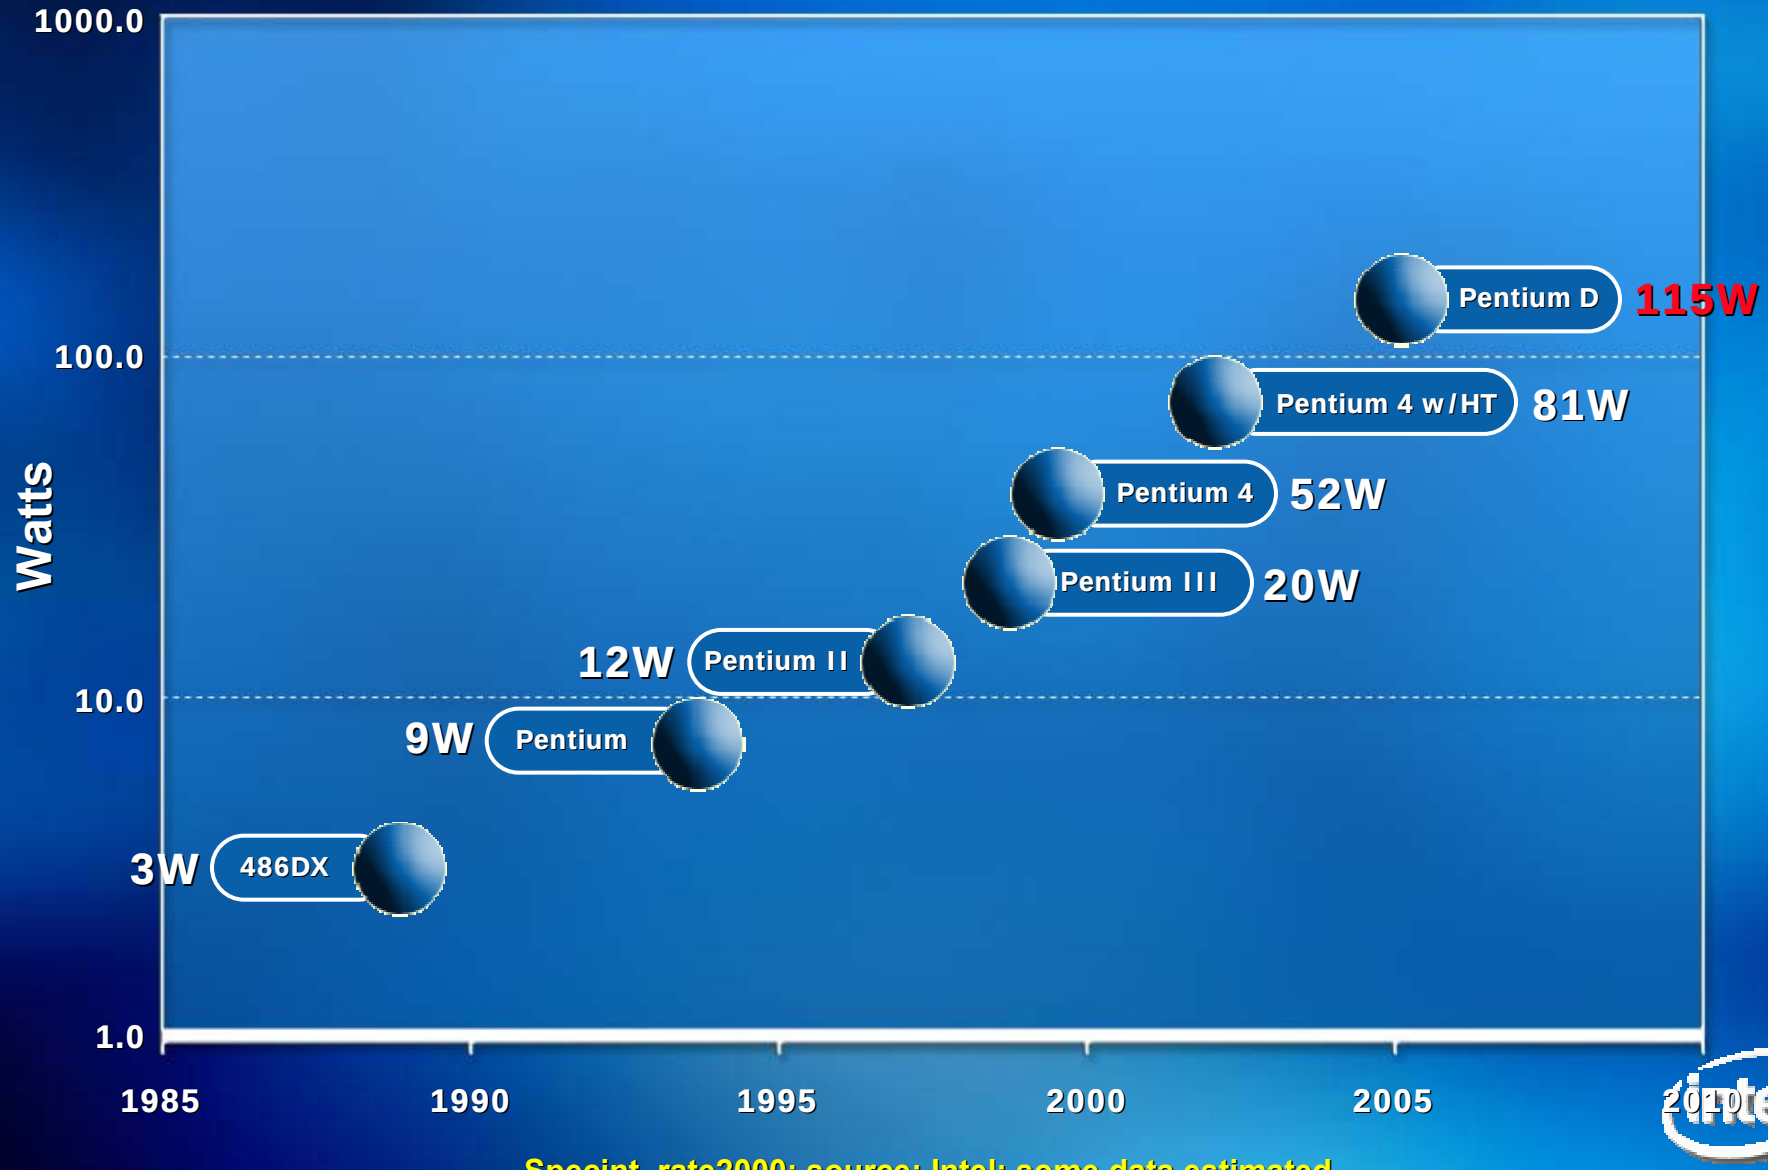

Specint_rate2000; source: Intel; some data estimated.

2003

Shift to **PERFORMANCE per WATT** in Notebooks

PC MAGAZINE
www.pcmag.com

March 12, 2003

"The new processor delivers on Intel's promise of faster performance and a battery-life boost measurable in hours."

intel inside
centrino™
MOBILE TECHNOLOGY

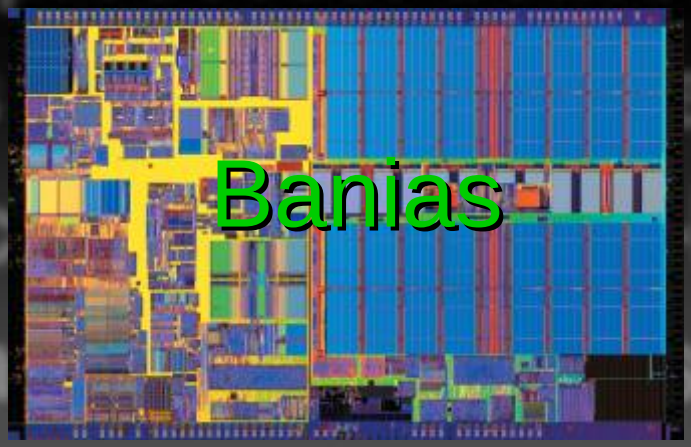

intel

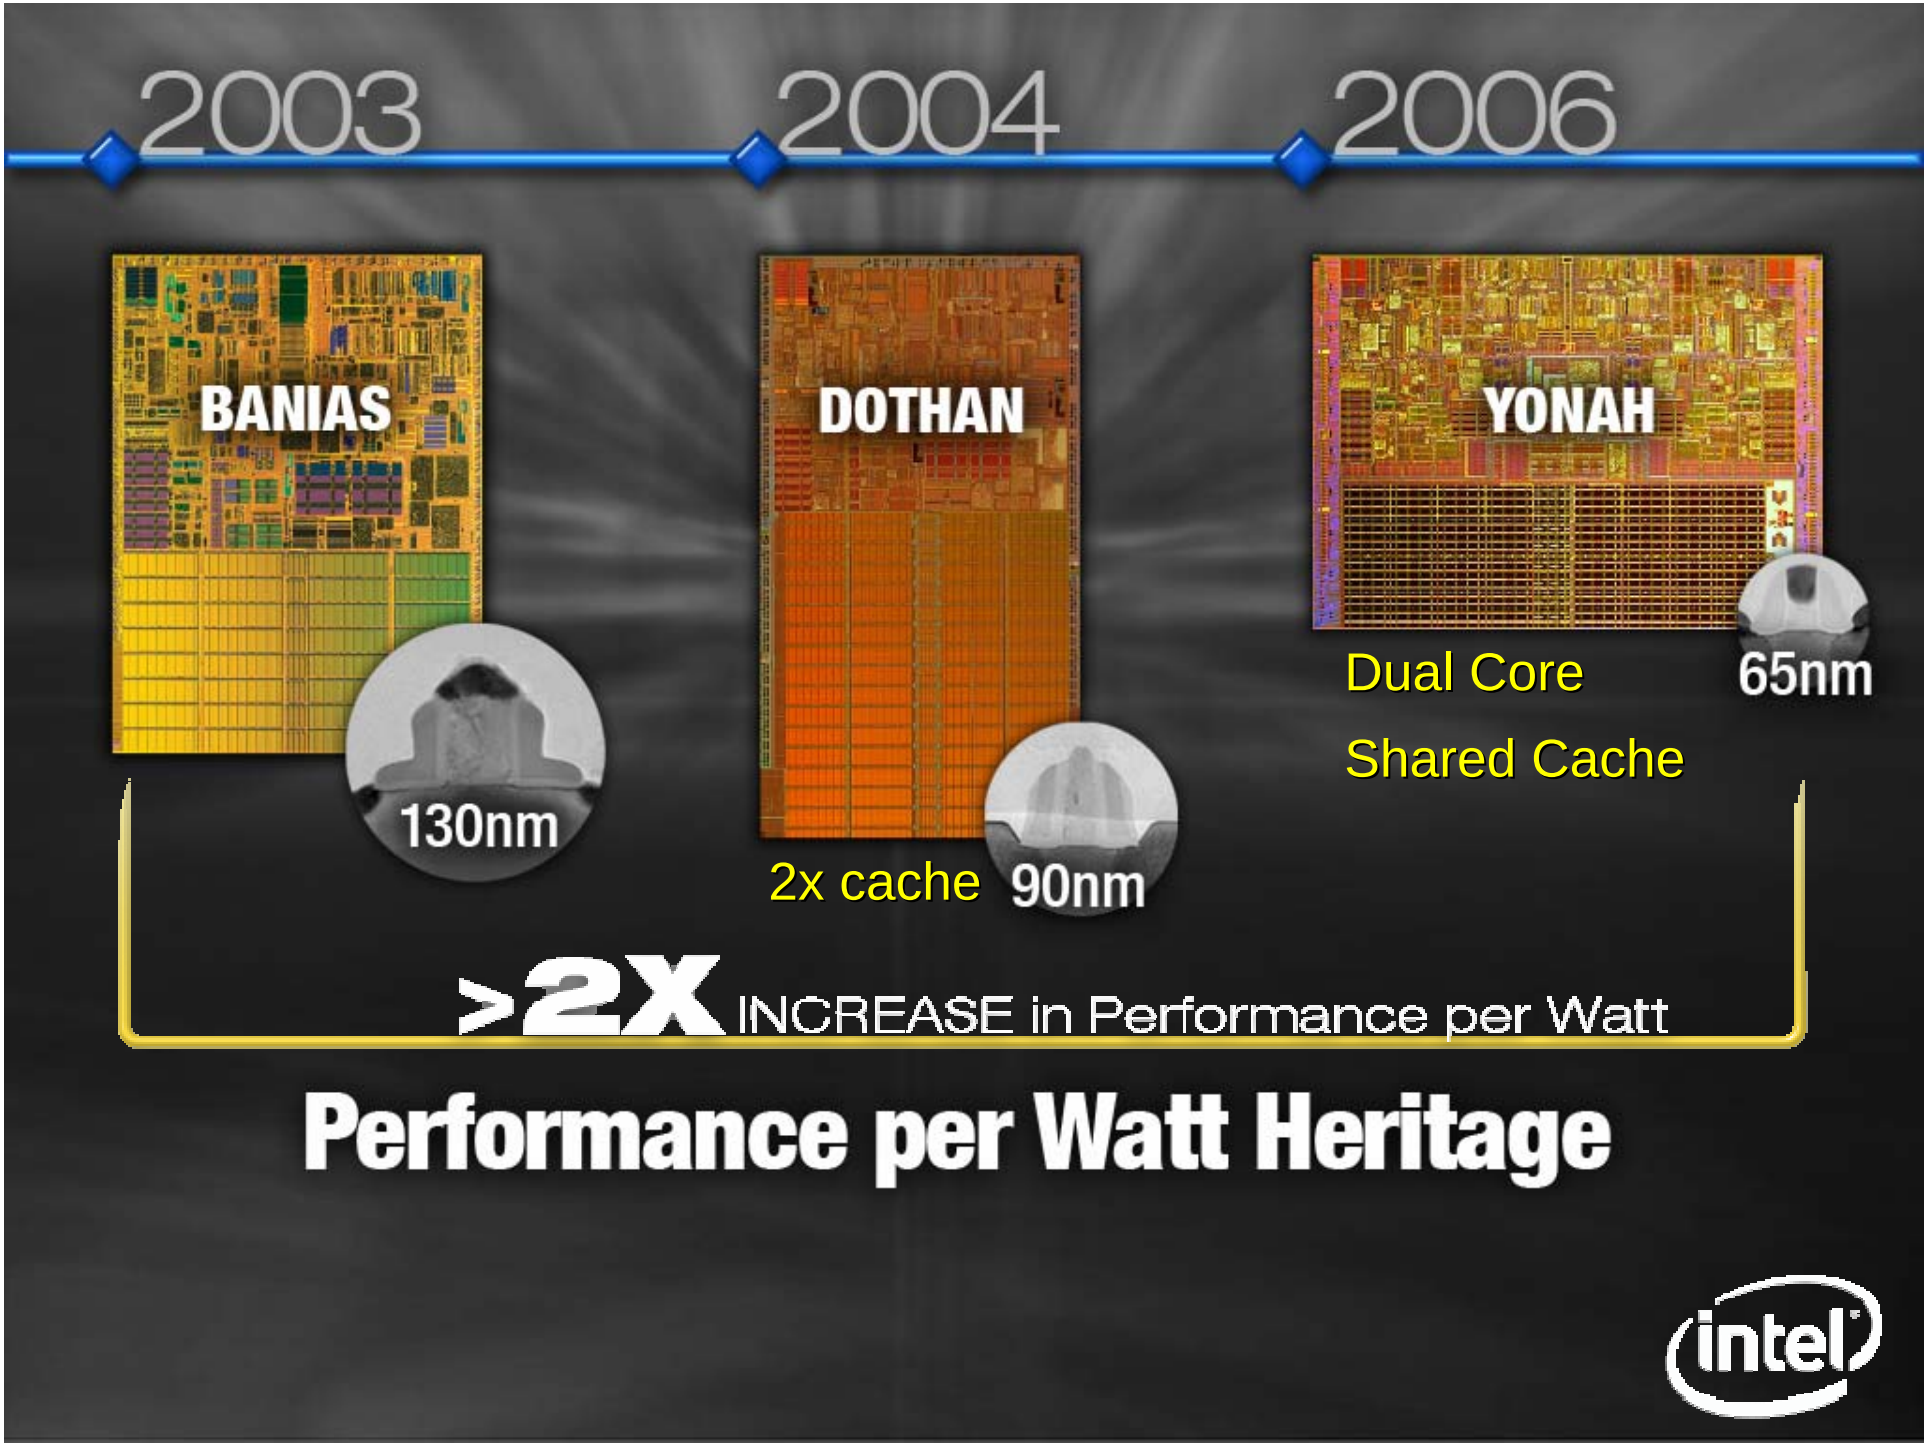
2003
2004
2006
BANIAS
130nm
DOTHAN
2x cache
90nm
YONAH
Dual Core
Shared Cache
65nm
>2X INCREASE in Performance per Watt
Performance per Watt Heritage
intel

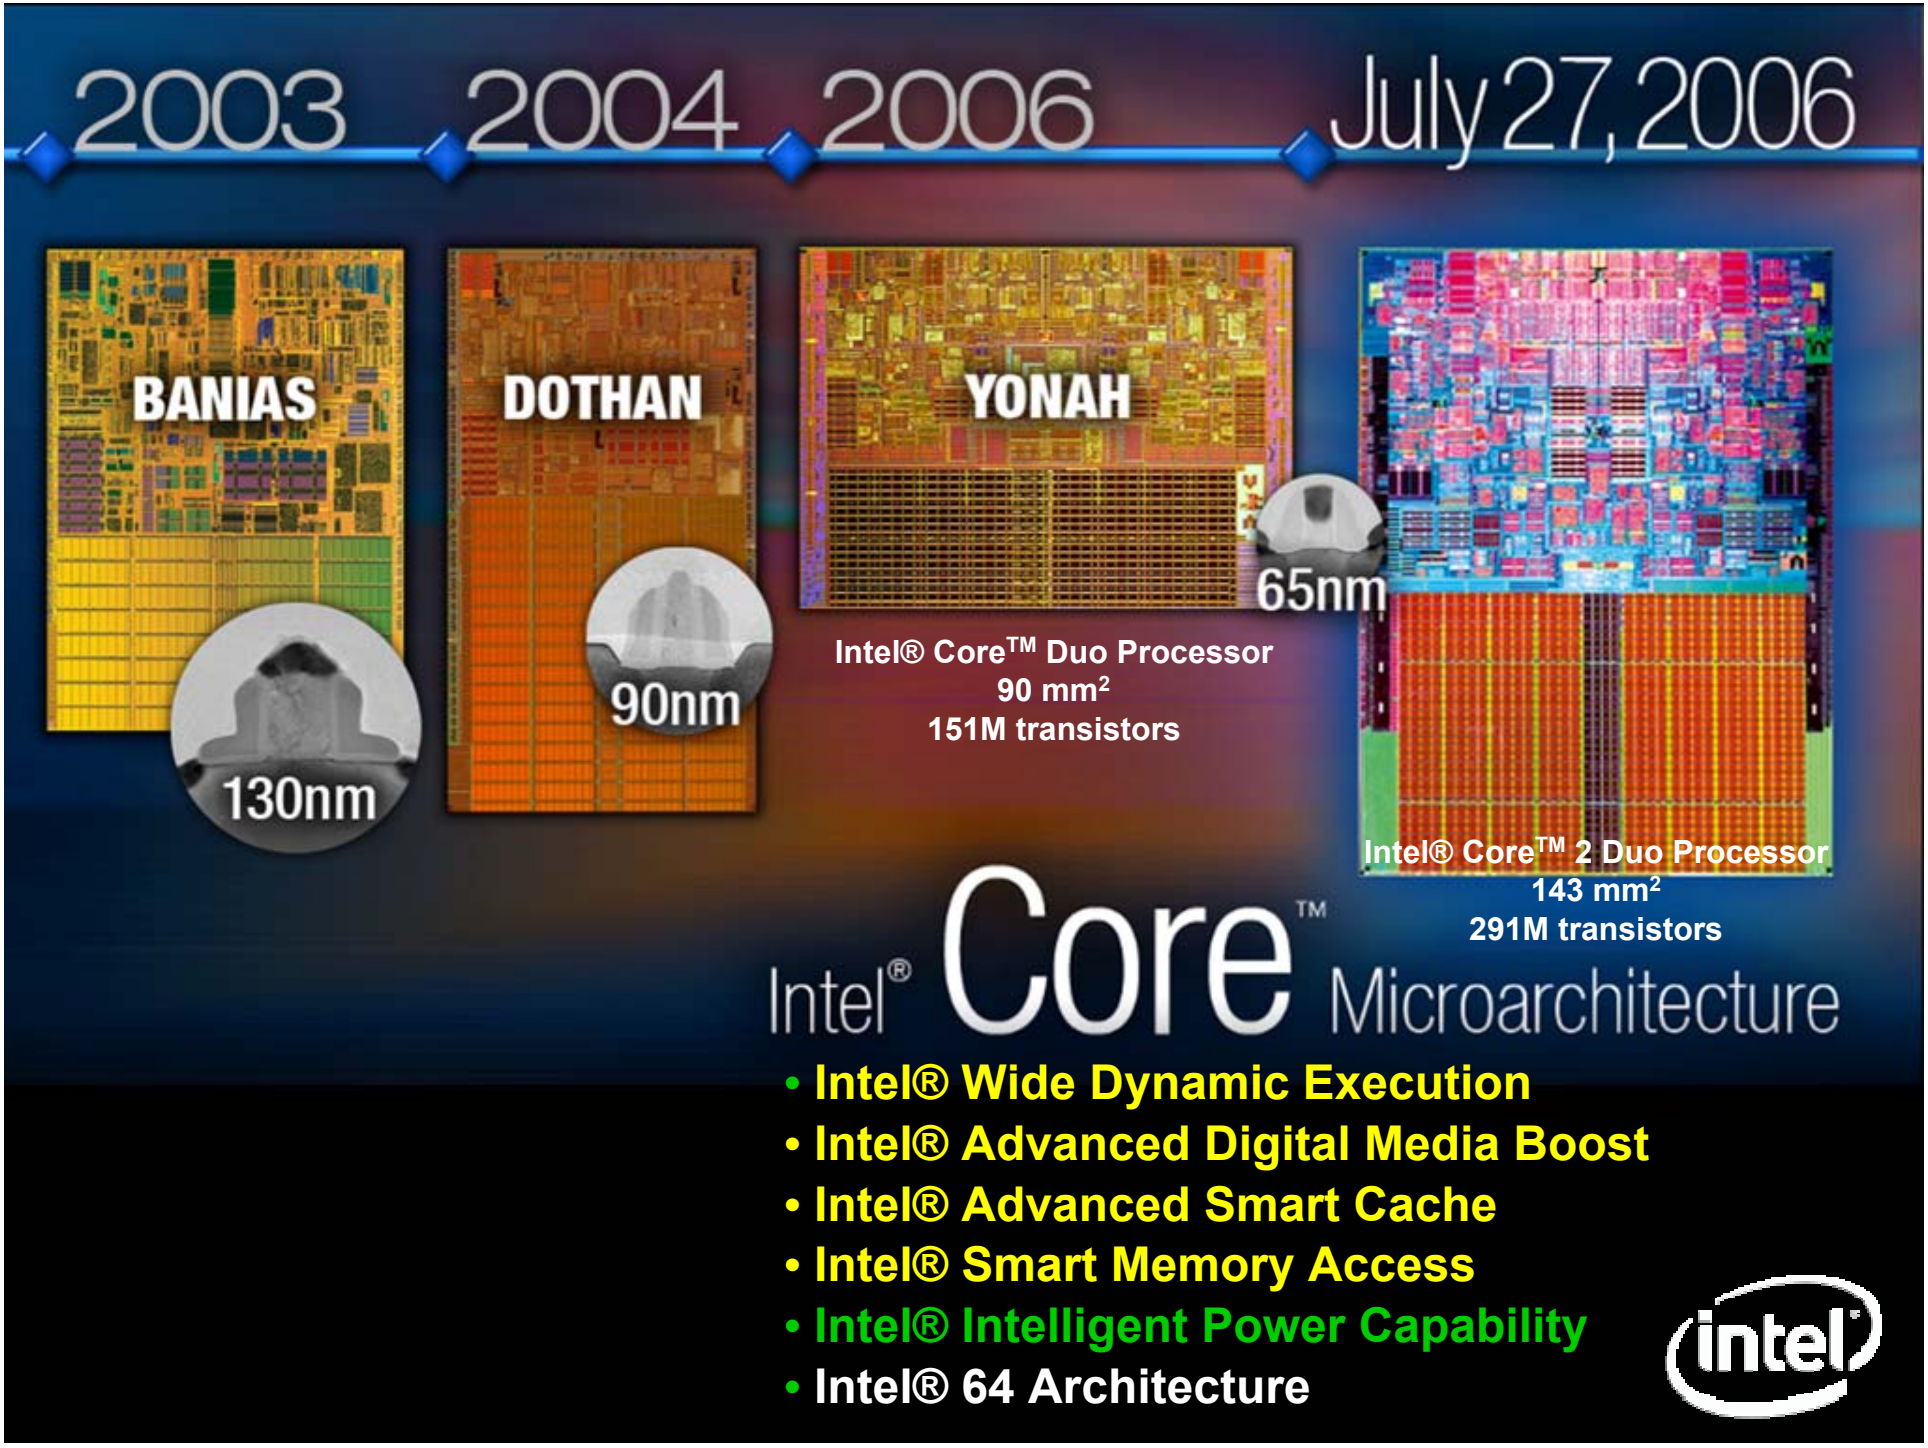
2003
2004
2006
July 27, 2006
BANIAS
130nm
DOTHAN
90nm
YONAH
65nm
Intel® Core™ Duo Processor
90 mm²
151M transistors
Intel® Core™ 2 Duo Processor
143 mm²
291M transistors
Intel® Core™ Microarchitecture
• Intel® Wide Dynamic Execution
• Intel® Advanced Digital Media Boost
• Intel® Advanced Smart Cache
• Intel® Smart Memory Access
• Intel® Intelligent Power Capability
• Intel® 64 Architecture
intel

Processor Architecture 101

Delivered Performance = Frequency * Instructions Per Cycle (IPC)

Goal is higher performance and lower power

Power α $C_{dynamic}$ * V * V * Frequency

$C_{dynamic}$ is roughly a product of area and activity
"how many bits" * "how much do they toggle"

Processor Architecture 101

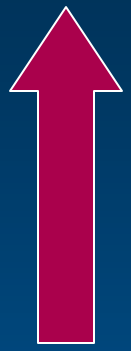

Delivered Performance = Frequency * Instructions Per Cycle (IPC)

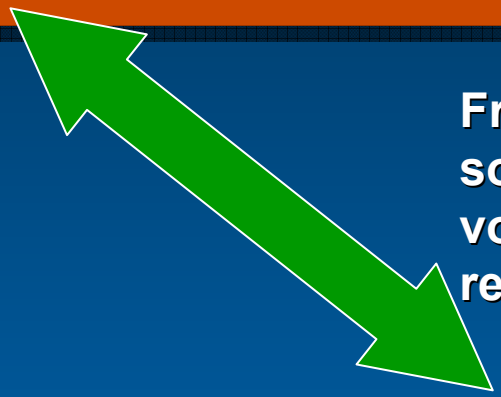

Frequency is proportional to voltage, so frequency reduction coupled with voltage reduction results in cubic reduction in power.

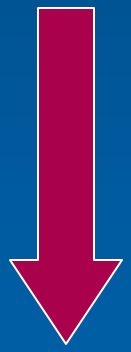

Power α $C_{dynamic}$ * V * V * Frequency

Processor Architecture 101

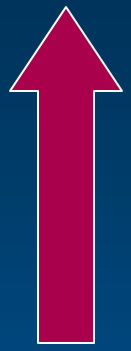

Delivered Performance = Frequency * Instructions Per Cycle (IPC)

Higher IPC usually results in wider data paths and/or more speculation : directly increasing C dynamic

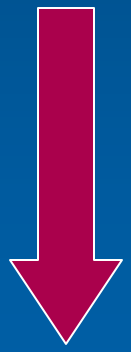

Power α $C_{dynamic}$ * V * V * Frequency

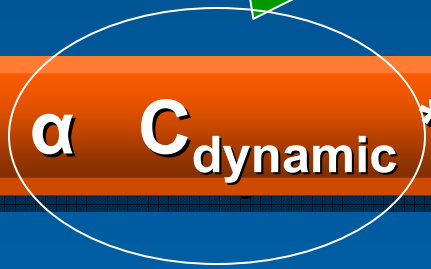

Micro-op Reduction

Delivered Performance = Frequency * Instructions Per Cycle (IPC)

Fewer uops per instruction allows IPC to be increased while lowering C dynamic (less bits and less toggling)

Power = $C_{dynamic}$ * V * V * Frequency

Techniques for Micro-op Reduction

- ESP Tracker (Extended Stack Pointer)
  - Execute Stack Pointer updates in dedicated hardware
  - ***Intel® Core™ microarchitecture increases BW by 33%****
- Micro-Op Micro-Fusion
  - Single Uop representation of "multi-uop" instruction
  - ***Intel® Core™ microarchitecture increase # instructions****
- ***Macro-Fusion***
  - New technique in Intel® Core™ microarchitecture

* Techniques pioneered on Intel® Pentium® M processors

Increasing Energy Efficiency

Specint_rate2000; source: Intel; some data estimated.

Intel® Intelligent Power Capability

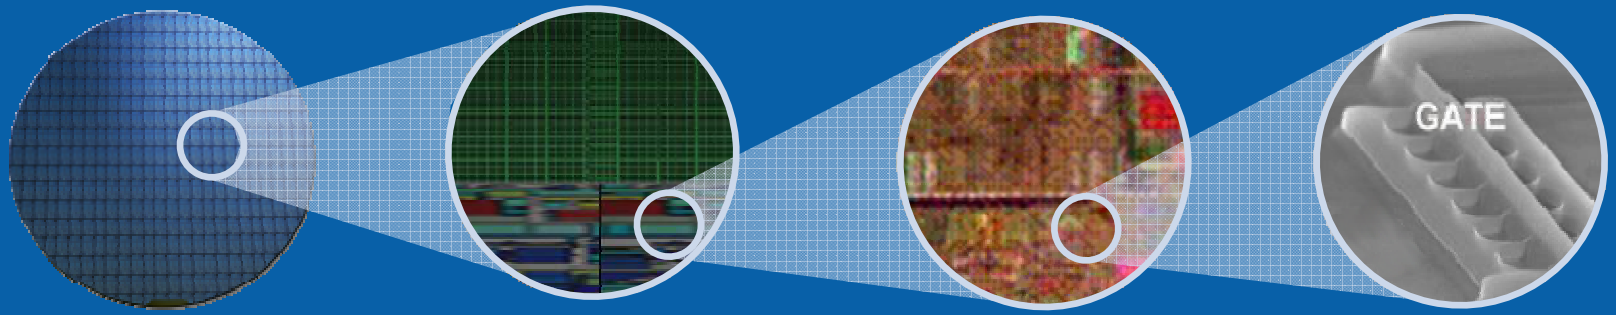

Process	Coarse Grained	Ultra Fine Grained	Transistor
• 65nm • Strained Silicon • Low-K Dielectric • More Metal Layers	• Aggressive Clock Gating • Enhanced Speed-Step	• Low VCC Arrays • Blocks Controlled Via Sleep Transistors	• Low Leakage Transistors • Sleep Transistors

Energy ADVANTAGE

- Mobile-Level Power Management
- Energy Efficient Performance

*Graphics not representative of actual die photo or relative size

Intelligent Power Capability

Extending the power management architecture

- Intel® Pentium® M processor innovated a new power management architecture
- Intel® Core™ Duo extended the Pentium® M processor capability to multi-core

New Power Features within each processor core

- Ultra fine-grained power control
- Split Busses
- Platformization of Power Management Architecture

Enhancing Energy Efficiency

Intelligent Power Capability

Ultra Fine Grained Power Control

Even during periods of high performance execution, many parts of the chip core can be shut off.

Example could be a SW memory initialization executing from front end with IQ operating as loop cache.

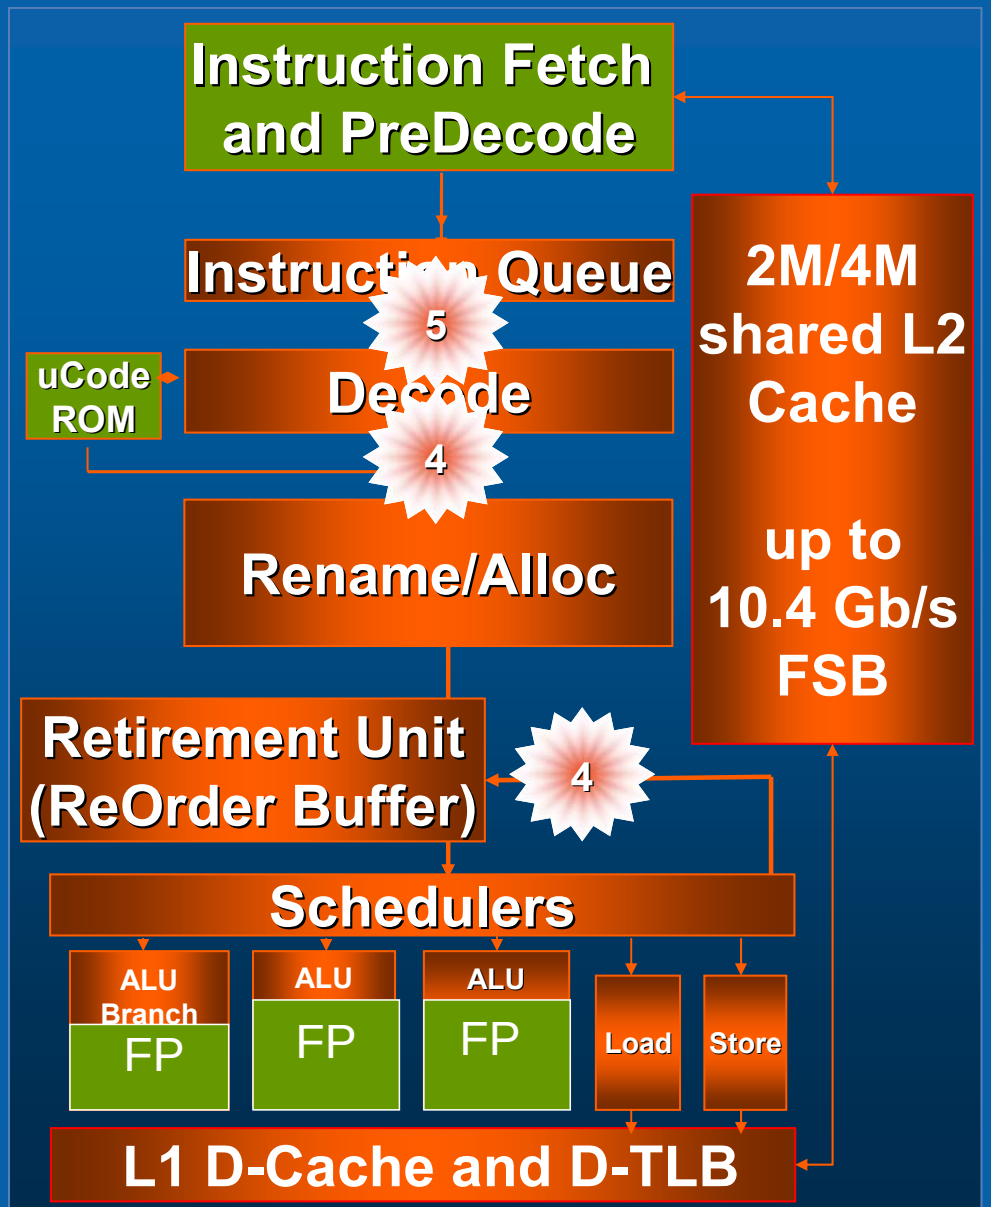

Intelligent Power Capability

Split Busses (core power feature)

Many buses are sized
for worst-case data

(x86 instruction of 15 bytes)
(ALU can write-back 128 bits)

Improved Energy Efficiency

Intelligent Power Capability

Split Busses (core power feature)

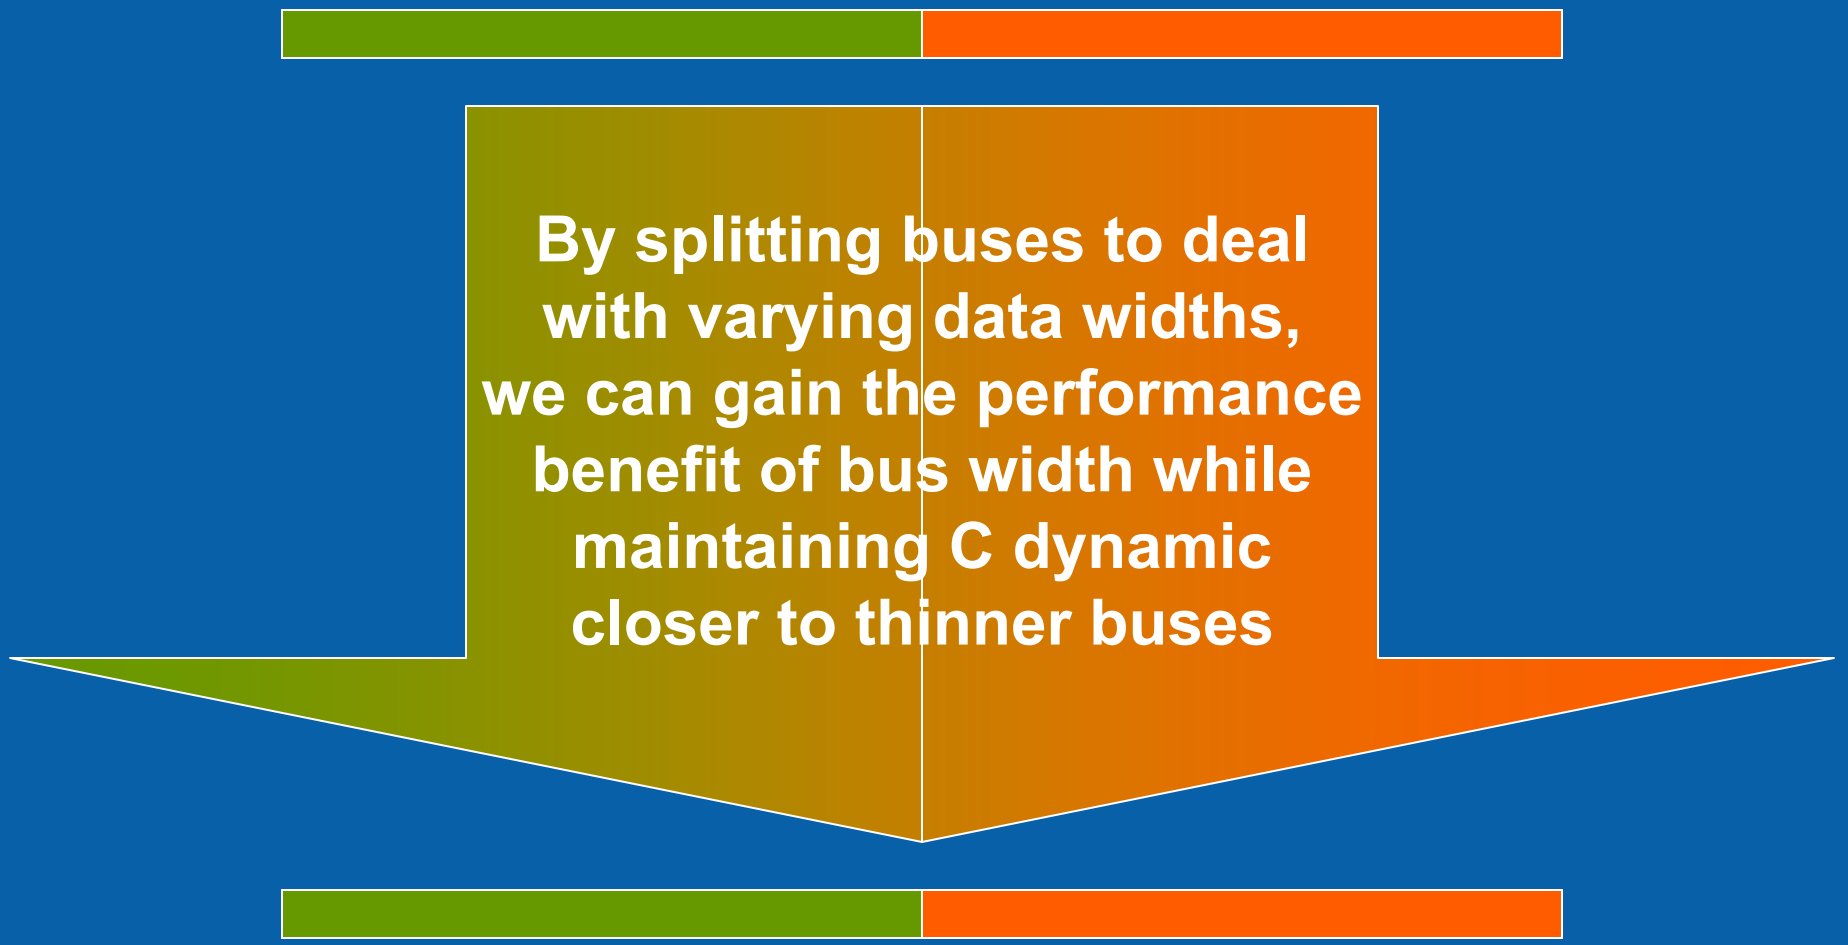

Improved Energy Efficiency

Platformization of Power Management Architecture

- Integrating best features from Server and Mobile products
- Exposing more to the system

- PSI-2 Power Status Indicator (Mobile)
- DTS Digital Thermal Sensors
- PECI Platform Environment Control Interface

Power Status Indicator (Mobile)

- Processor communicates power consumption to external platform components
  - –Optimization of voltage regulator efficiency
  - –Load line and power delivery efficiency

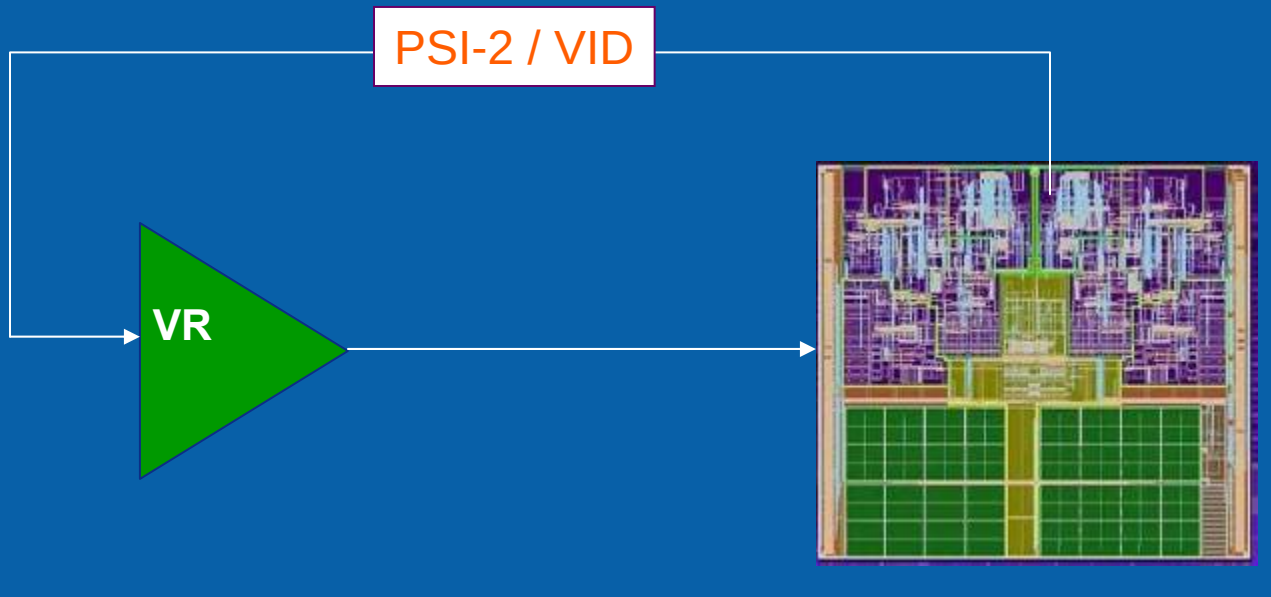

Enabling Efficient Processor and Platform Thermal Control…

DTS – Digital Thermal Sensor

- Several thermal sensors are located within the Processor to cover all possible hot spots
- Dedicated logic scans the thermal sensors and measures the maximum temperature on the die at any given time
- Accurately reporting Processor temperature enables advanced thermal control schemes

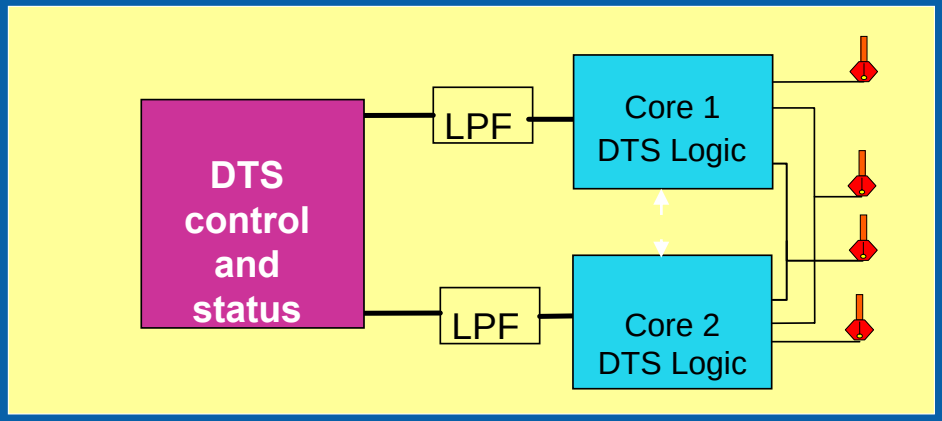

Platform Environment Control Interface (PECI)

- Processor provides its temperature reading over a **multi drop single wire bus** allowing efficient platform thermal control

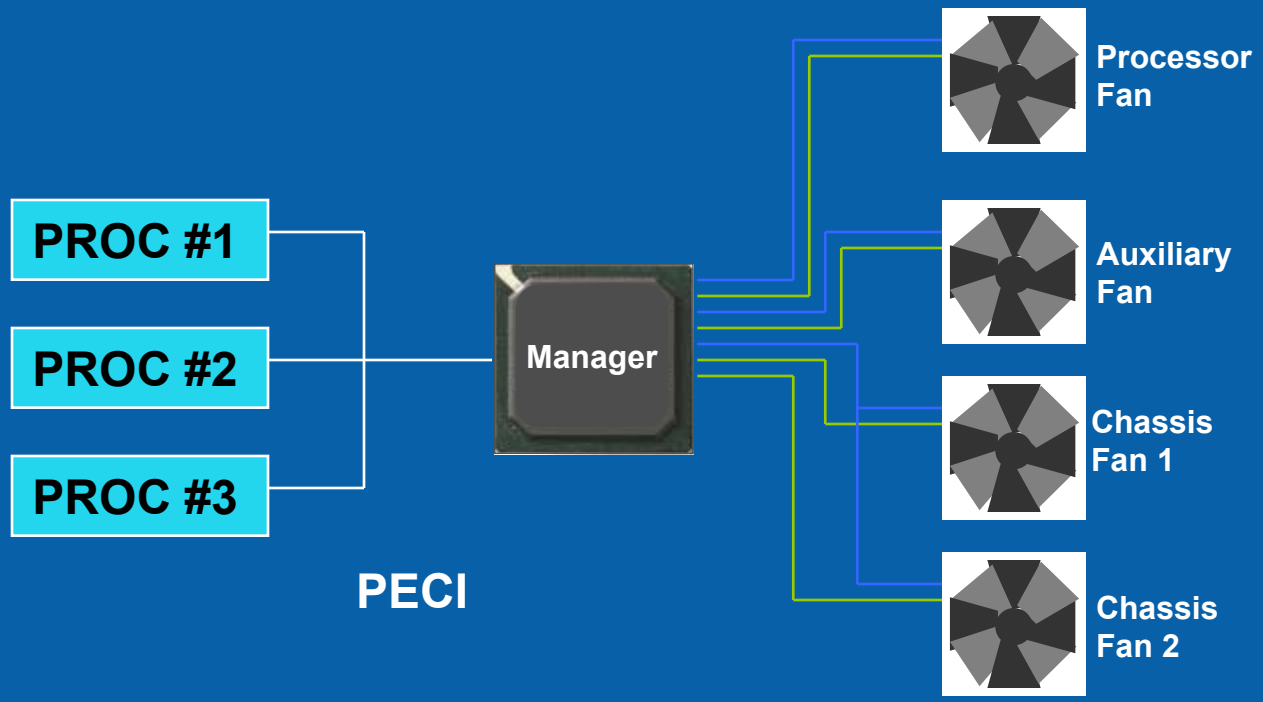

Intel® Core™ Microarchitecture Products

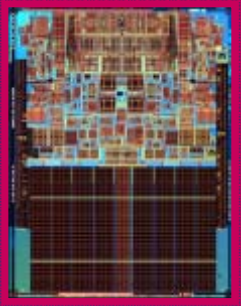

Higher Performance at Lower Power!

Five Major Innovations

Spans Mobile, Desktop & Server

Core™ 2 Duo *Mobile*

Perf >20%

Energy Equal

Relative to Intel® Core Duo™ T2600

Core™ 2 Duo *Desktop*

Perf +40%

Energy - 40%

Relative to Pentium® D 950

Xeon® 5100 *DP Server*

Perf +80%

Energy - 35%

Relative to Xeon® 2.8GHz 2x2MB

Energy Efficient Performance Leadership

Source: Intel based on estimated SPECint*_rate_base2000 and thermal design power

Managing Power and Cooling Efficiency

Silicon:
Moore's law, Strained silicon, Transistor leakage control techniques, Clock gating

Processor:
Power scaling with load,
Policy based power allocation
Multi-threaded cores

System Power Delivery:
Fine grain power management,
Ultra fine grain power management

Facilities:
Air cooling and liquid cooling options
Vertical integration of cooling solutions

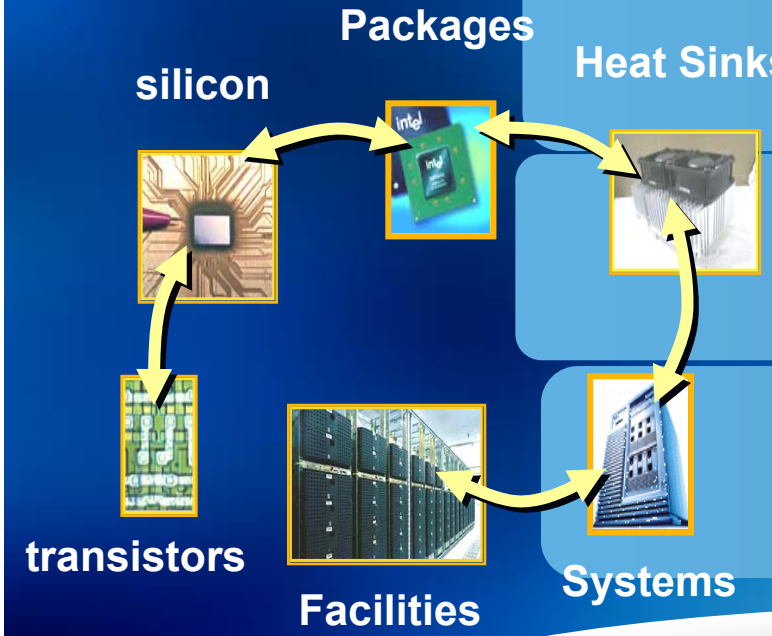

Power management

from transistors to facilities

Transition from Classical Scaling to Power-Efficient Scaling

1970's: Classical/Constant-Vdd -> PMOS, NMOS

1980's: Classical/Constant Vdd -> CMOS

1990's: Classical/Voltage scaling -> ($P = CV^2f$)

2000's: Power-efficient scaling -> New processes, materials, structures; co-optimization with design

Increased Transistor Performance With 2nd Generation Strained Silicon

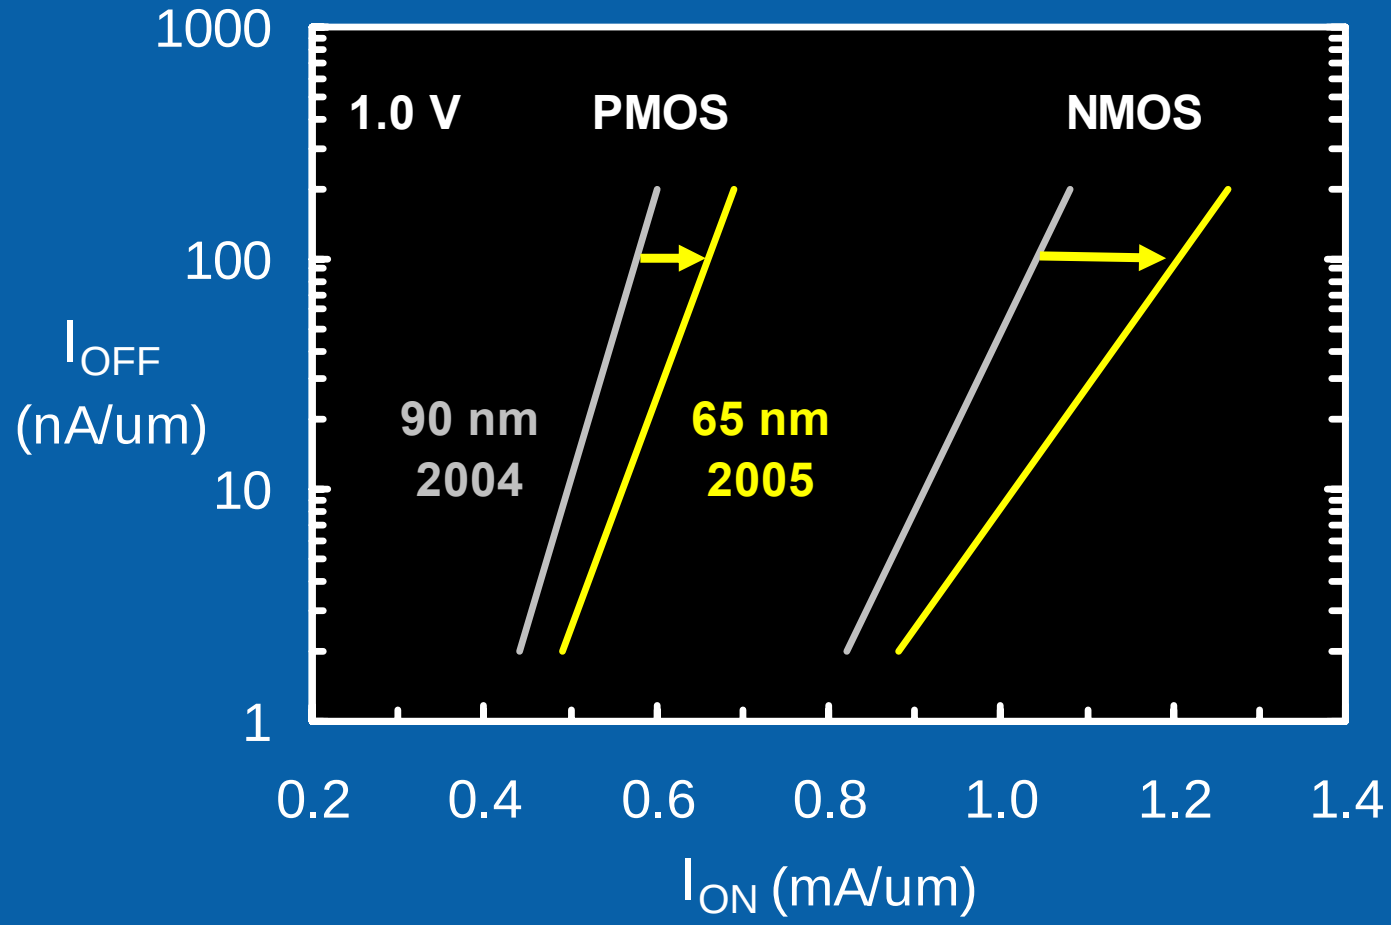

~16% higher drive current plus ~20% lower gate capacitance improves transistor switching speed by >20% at same leakage level

Reduced Transistor Leakage With 2nd Generation Strained Silicon

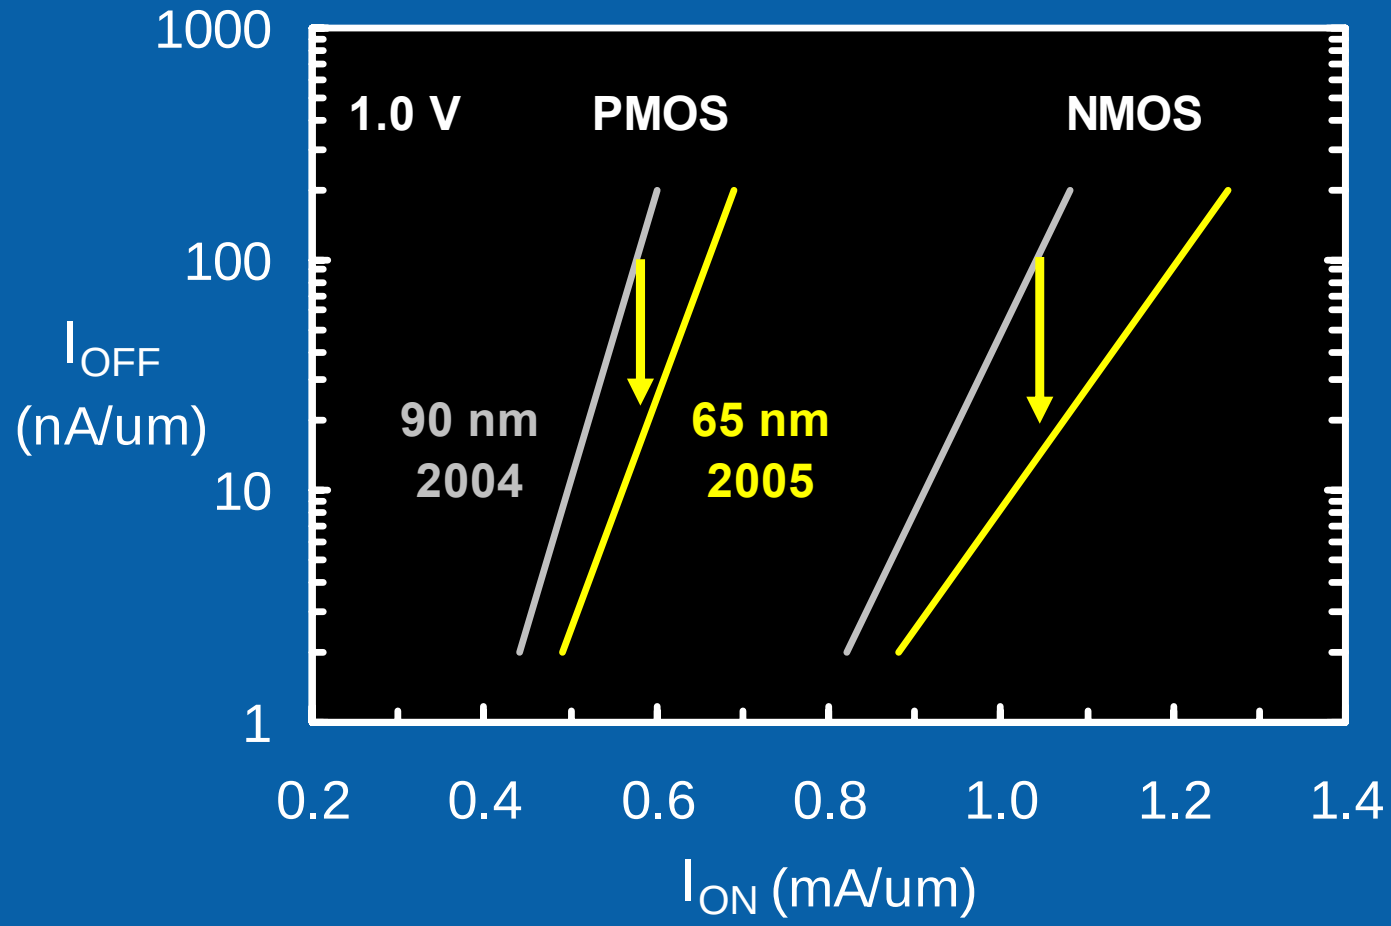

Or, transistor leakage can be reduced ~5x at same drive current

Power Efficient Interconnect Includes 2nd Generation CDO

- Low-k carbon doped oxide (CDO) dielectric reduces interconnect capacitance (improved from 90 nm generation)
- Metal 8 layer is added for improved density and performance (1 more layer than 90 nm generation)
- Interconnect capacitance is reduced by use of low-k dielectric and by ~0.7x line length scaling
- Lower capacitance improves interconnect performance and reduces chip power

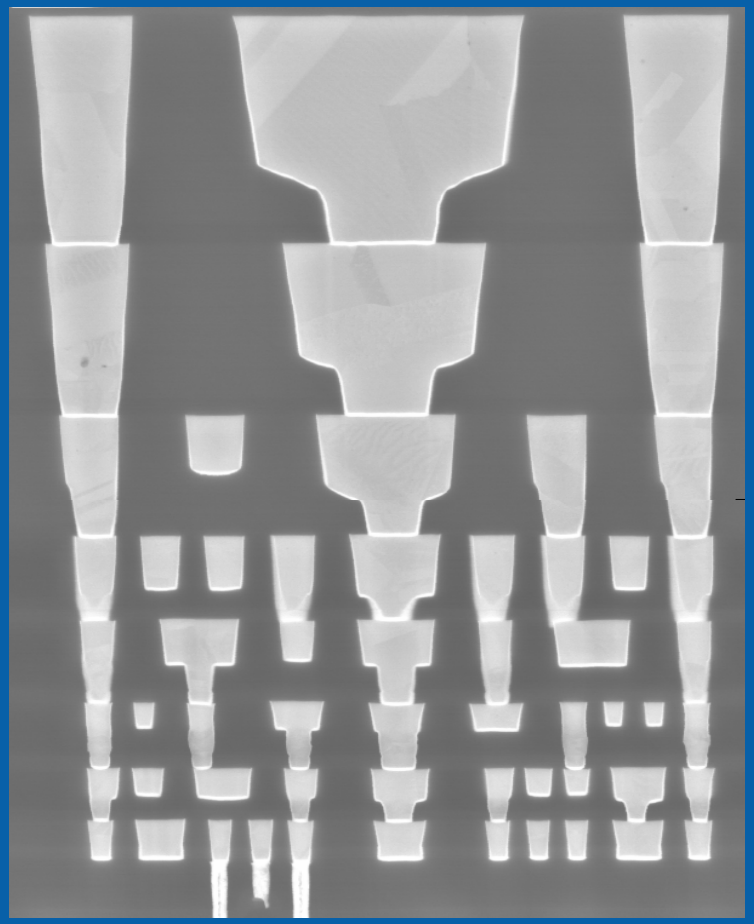

Sleep Transistors Reduce Leakage Power

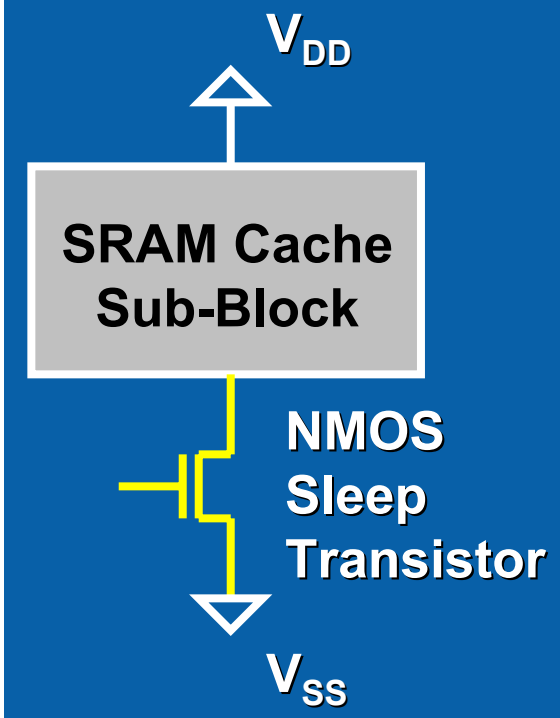

70 Mbit SRAM IR photos

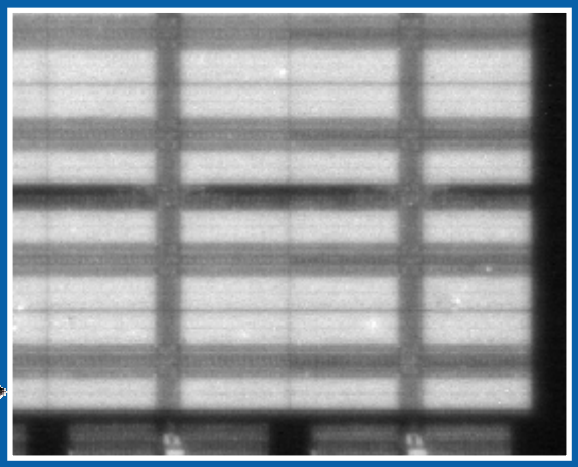

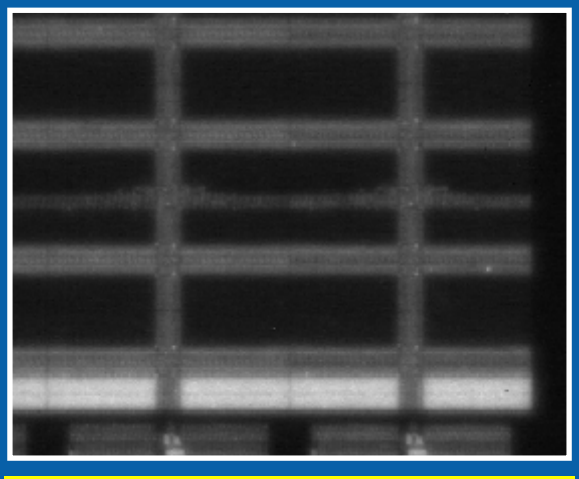

Normal SRAM sub-block leakage

Sleep transistors shut off leakage in inactive sub-blocks

>3x SRAM leakage reduction with use of sleep transistors

45 nm Logic Process on Track for Delivery in 2007

Process Name	P1262	P1264	P1266	P1268
Lithography	90 nm	65 nm	45 nm	32 nm
1st Production	2003	2005	2007	2009

Moore's Law continues!

Intel continues to develop a new technology generation every 2 years

45 nm Technology Benefits

Compared to today's 65 nm technology, the 45 nm technology will provide the following product benefits:

- ~2x improvement in transistor density, for either smaller chip size or increased transistor count
- >20% improvement in transistor switching speed or >5x reduction in leakage power
- >30% reduction in transistor switching power

This process technology will provide the foundation to deliver improved performance/Watt

Gate Dielectric Today is Only a Few Molecular Layers Thick

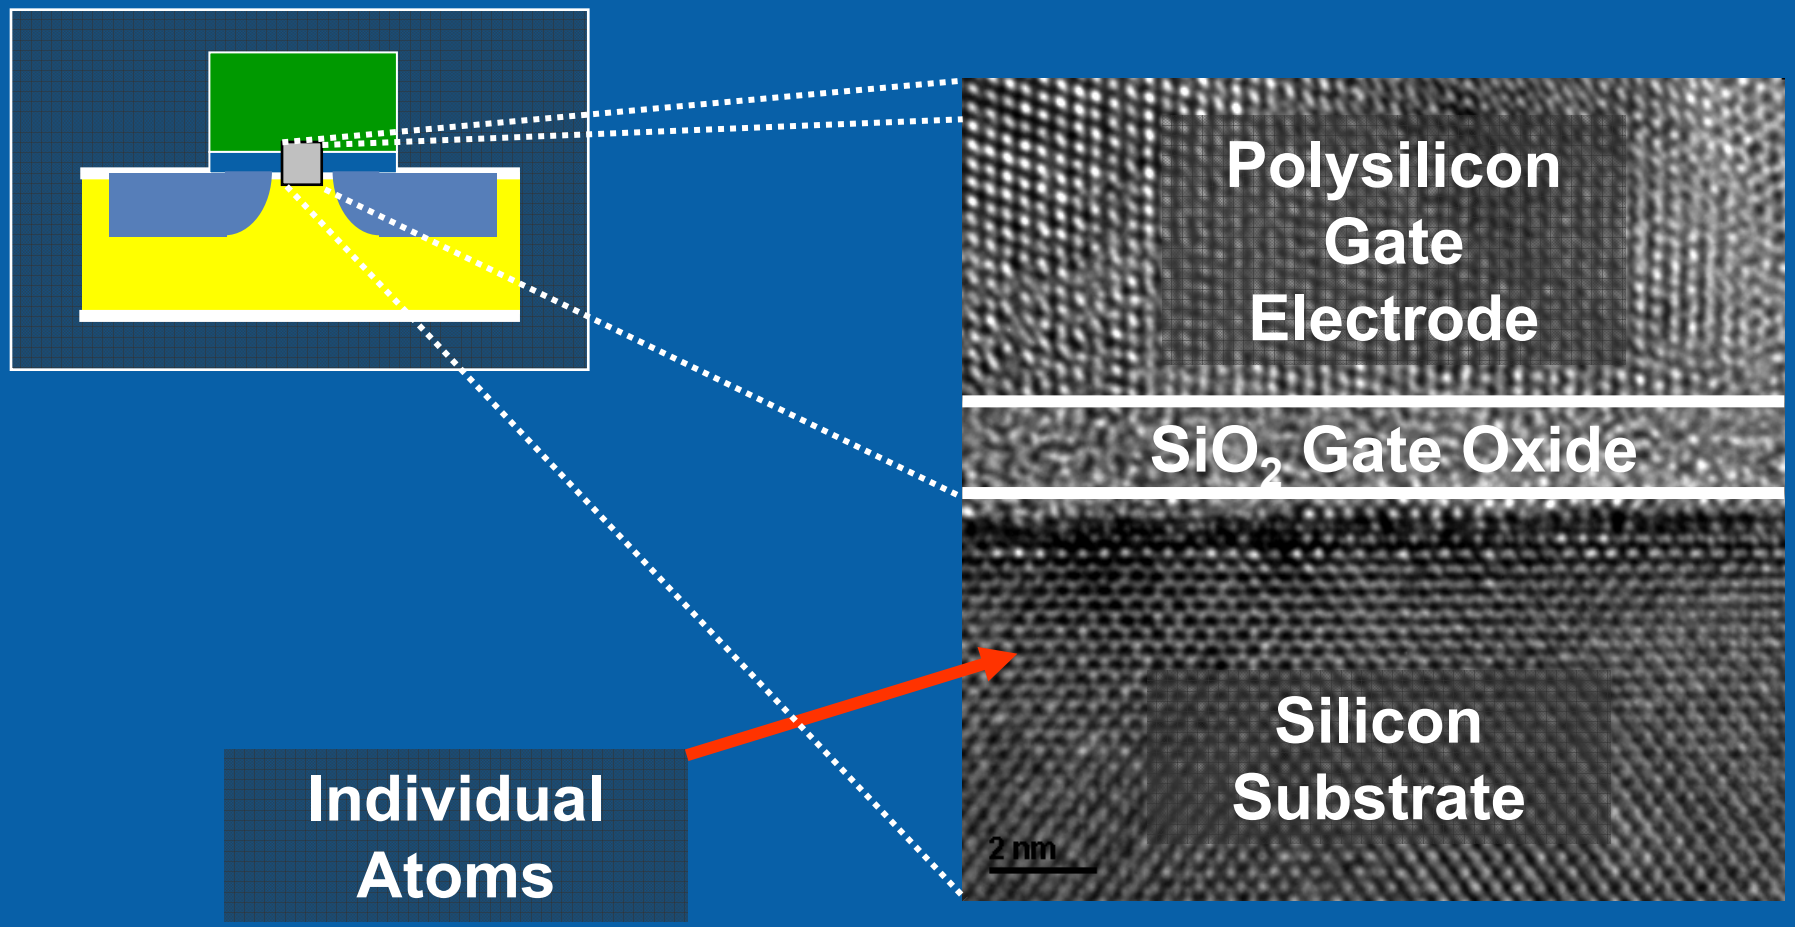

High-k Demonstrated >100x Leakage Reduction

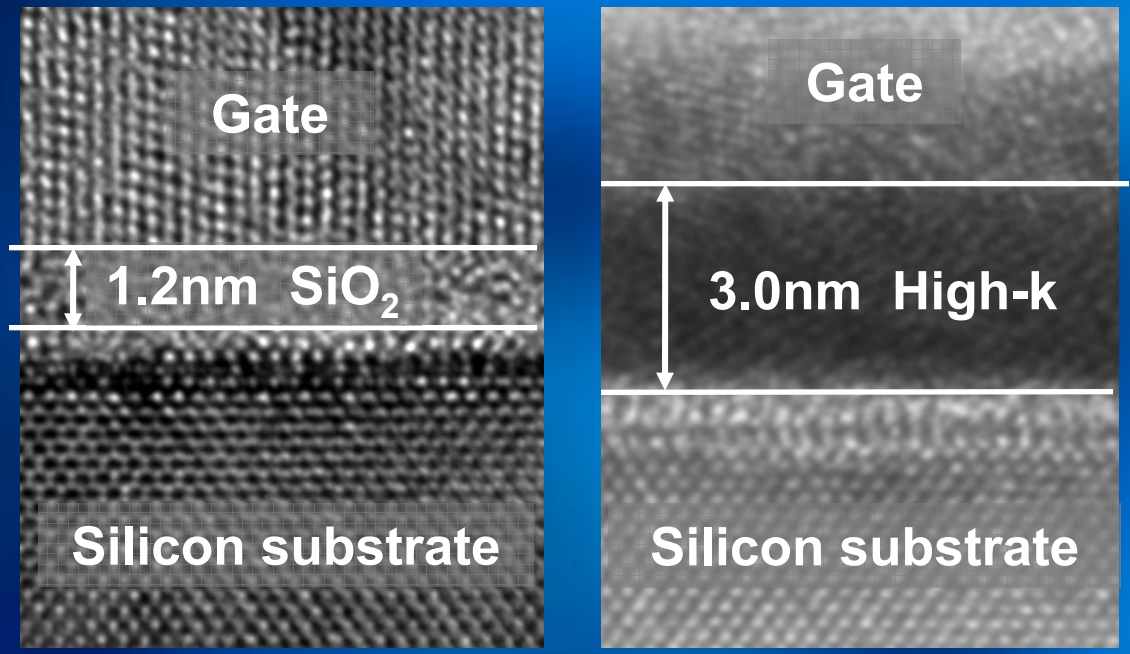

Benefits compared to current process technologies

	High-k vs. SiO_2	Benefit
Capacitance	60% greater	*Much faster transistors*
Gate dielectric leakage	> 100x reduction	*Far cooler*

Leakage Control

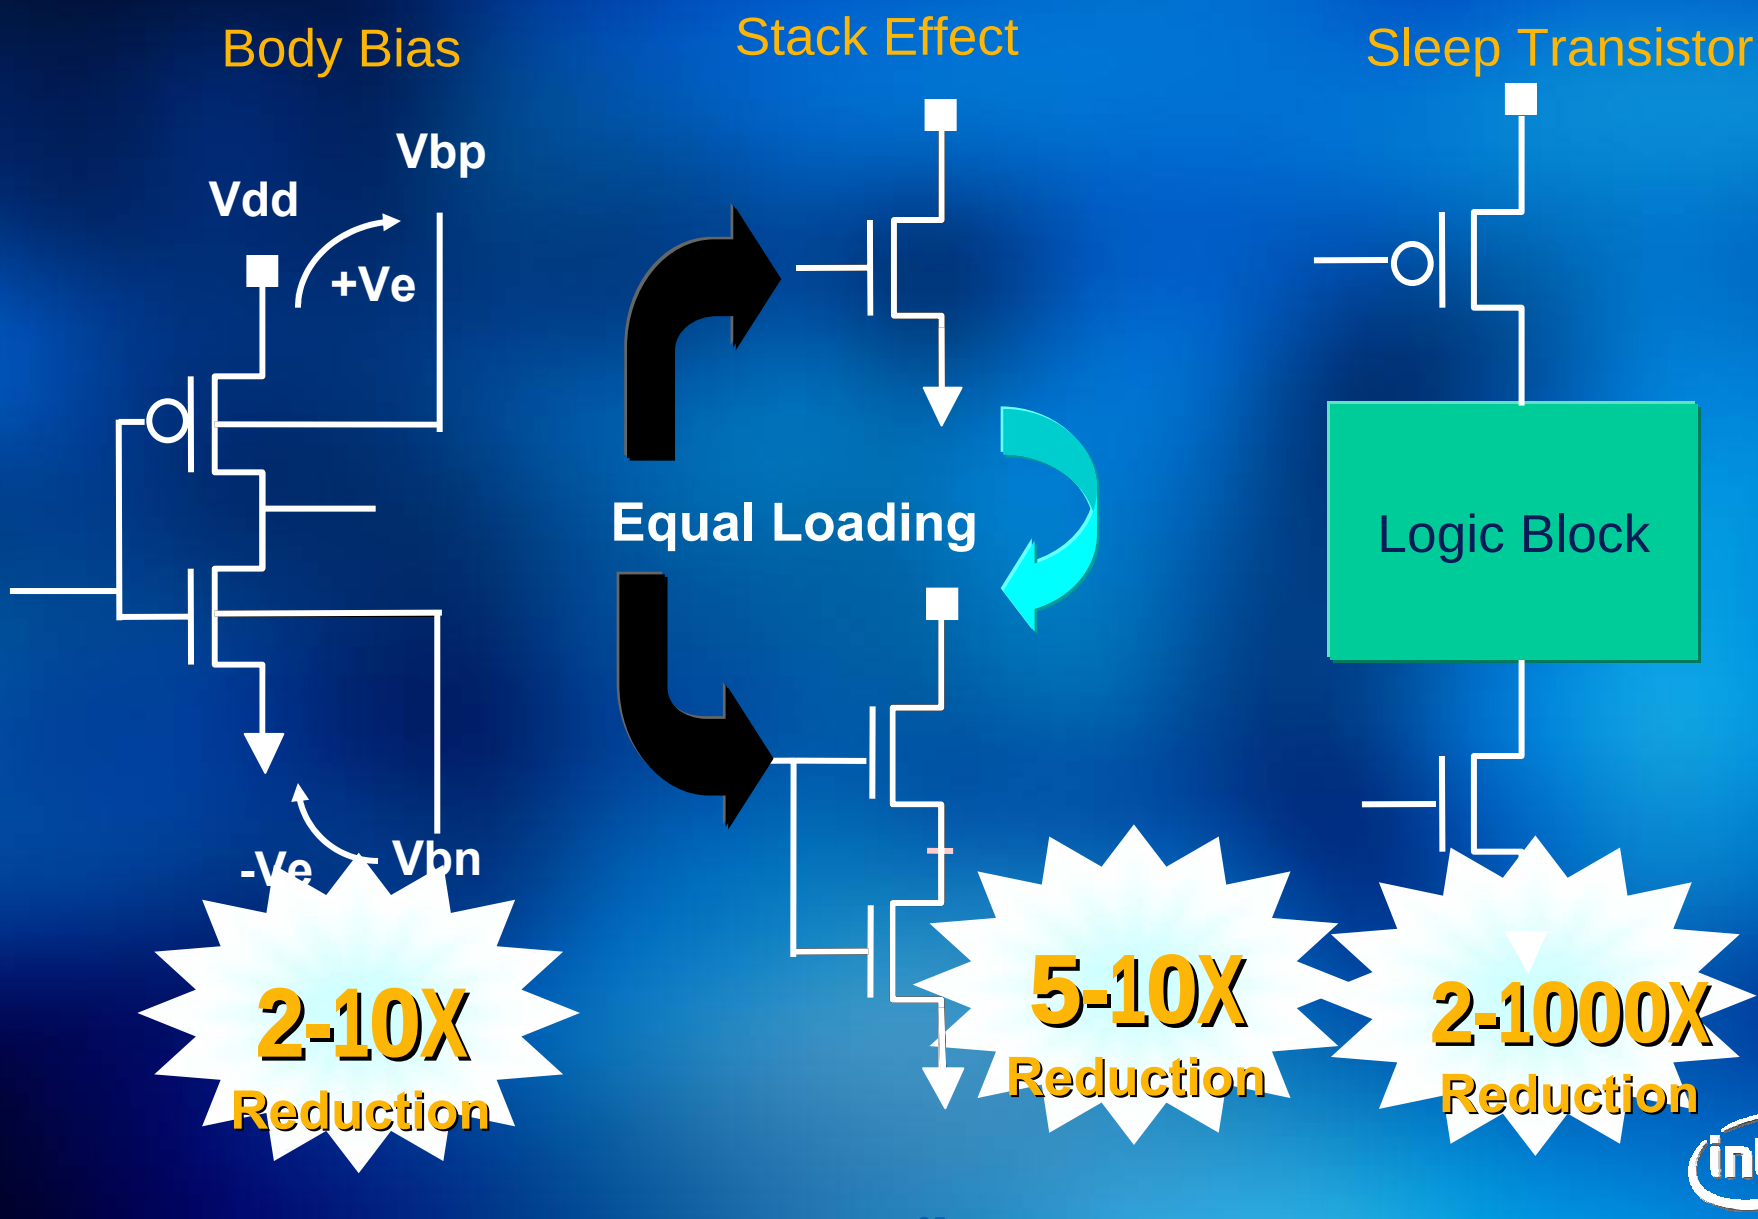

Active Power Reduction

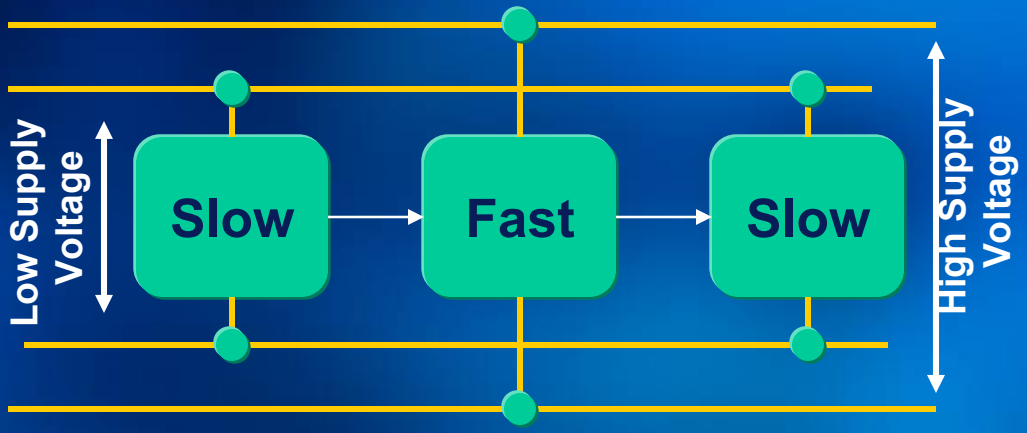

Multiple Supply Voltages

Throughput Oriented Designs

Vdd

Logic Block

Freq = 1
Vdd = 1
Throughput = 1
Power = 1
Area = 1
Pwr Den = 1

Vdd/2

Logic Block

Logic Block

Freq = 0.5
Vdd = 0.5
Throughput = 1
Power = 0.25
Area = 2
Pwr Den = 0.125

µArchitecture Techniques

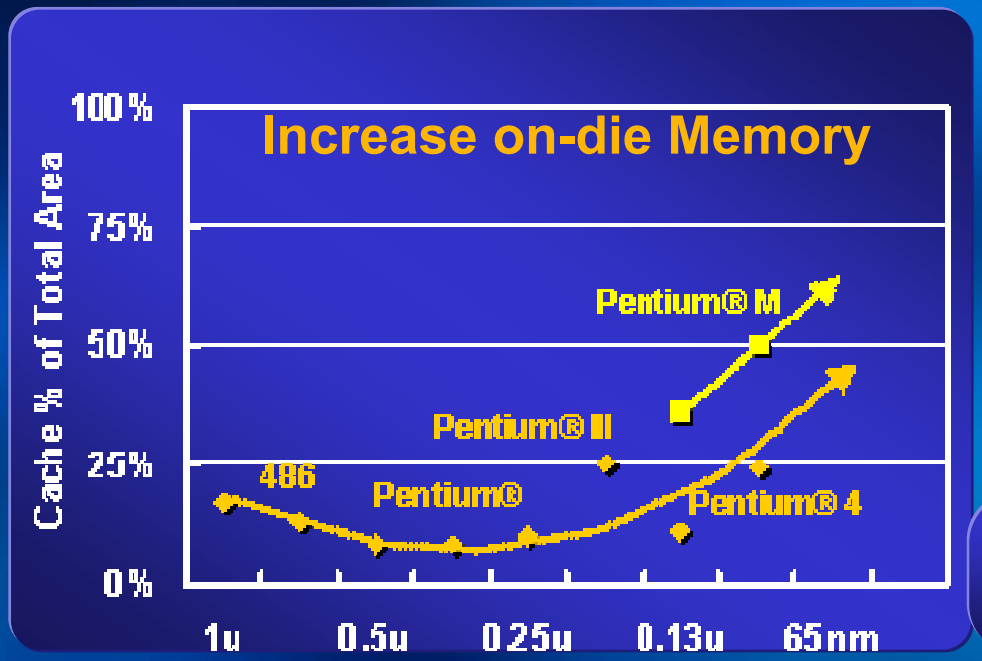

Multi-threading

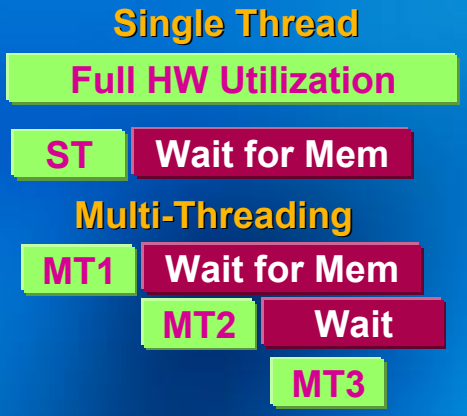

Improved performance, no impact on thermals & power delivery

Chip Multi-processing

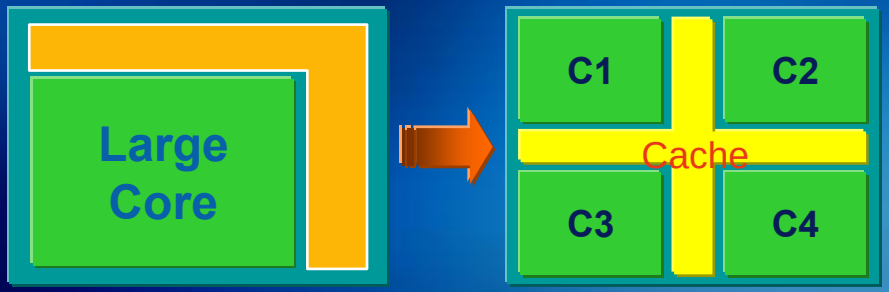

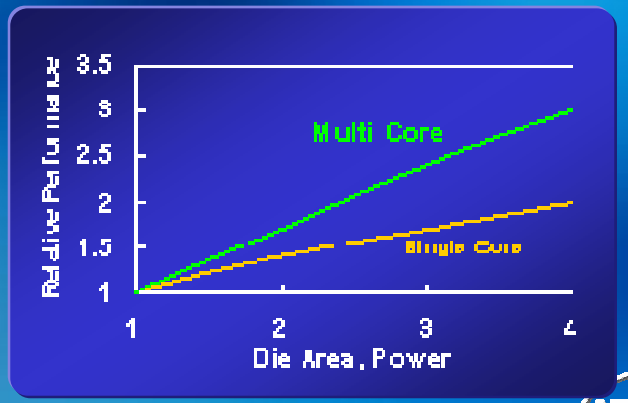

Energy Efficient Performance

$$\text{Performance Per Watt} = \frac{\text{System Performance}}{\text{System Power Consumption}}$$

- Customers want to know how much power the system is consuming under real workloads – “Watts at the Wall”
- Use external power measurement tool
- System connects to Wall power outlet via Power Meter

Breakdown of Platform Power

Desktop

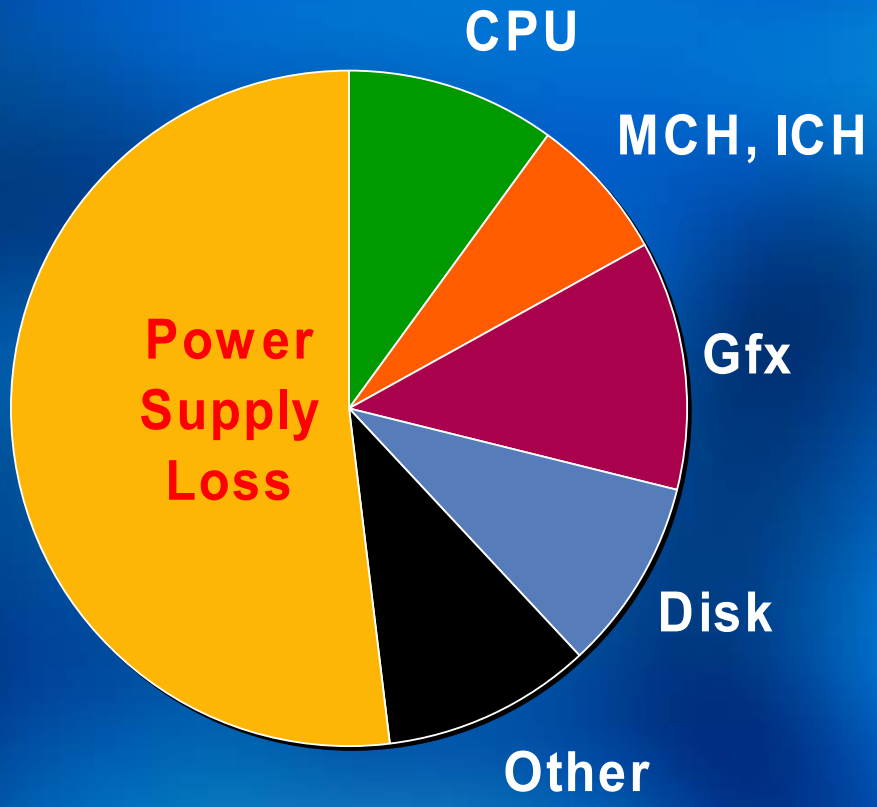

CPU & electronics power in a platform is low
Must reduce power of other platform ingredients

Power Delivery Improvements

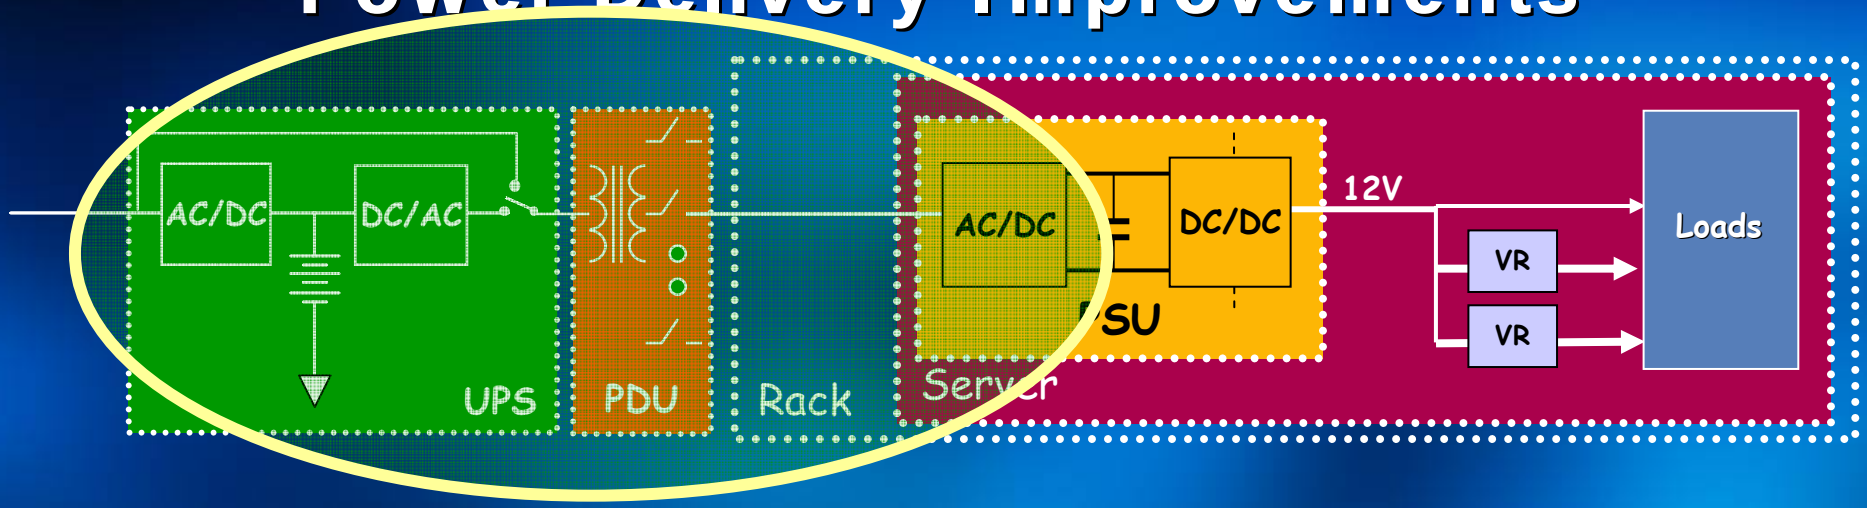

Data Center Power Delivery

Significant improvement in efficiency feasible

- Conversion stages can be eliminated
- Cooling benefits can be achieved

Power Delivery Improvements

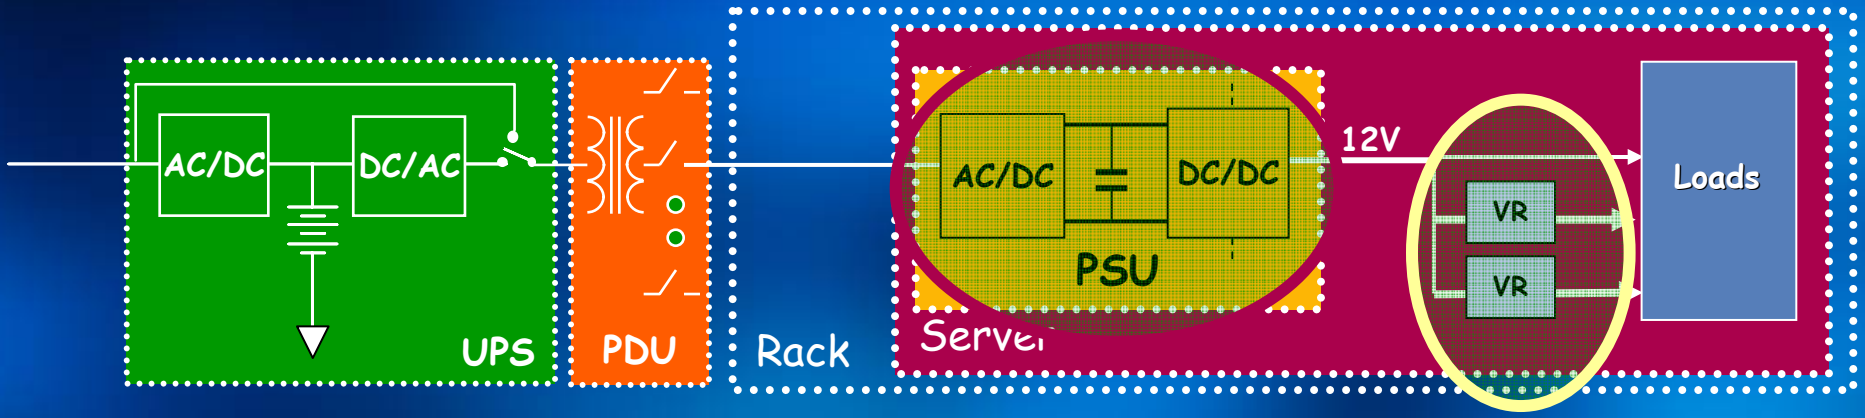

VR

- Efficiency continually increasing
- Phase dropping at light loads

PSU

- Typical PSU efficiency is 75%
- 80+% efficiency programs
  - EPA Energy Star (proposed), 80PLUS program, SSI specifications
- Technology exists for ~90% efficient PSUs
  - Cost remains a barrier to use, despite recuperation of cost within 1 year

Pathfinding for 90+% PSU
Light Load (5%) to Full Load (100%)

Load Adaptive PSU

- Extend platform power management scope to include power supply
- Extend multi-phase concept to DCX and AC/DC
  - To improve light load efficiency

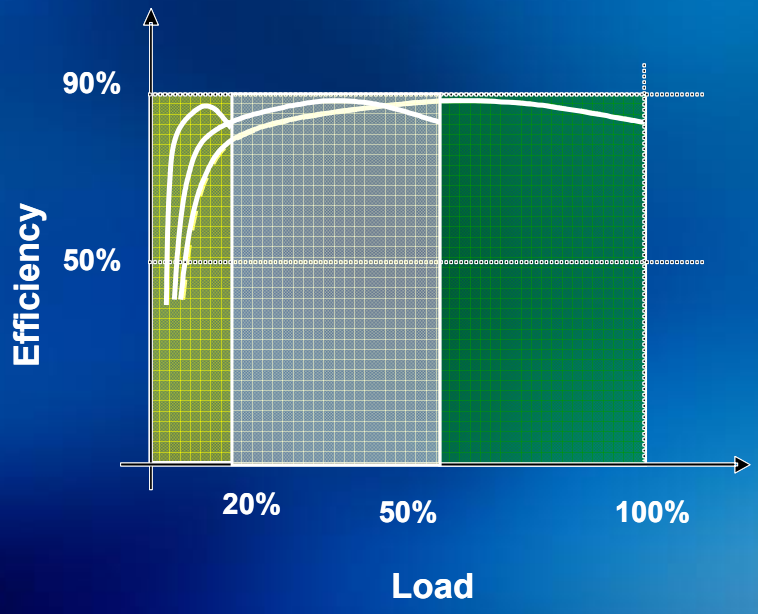

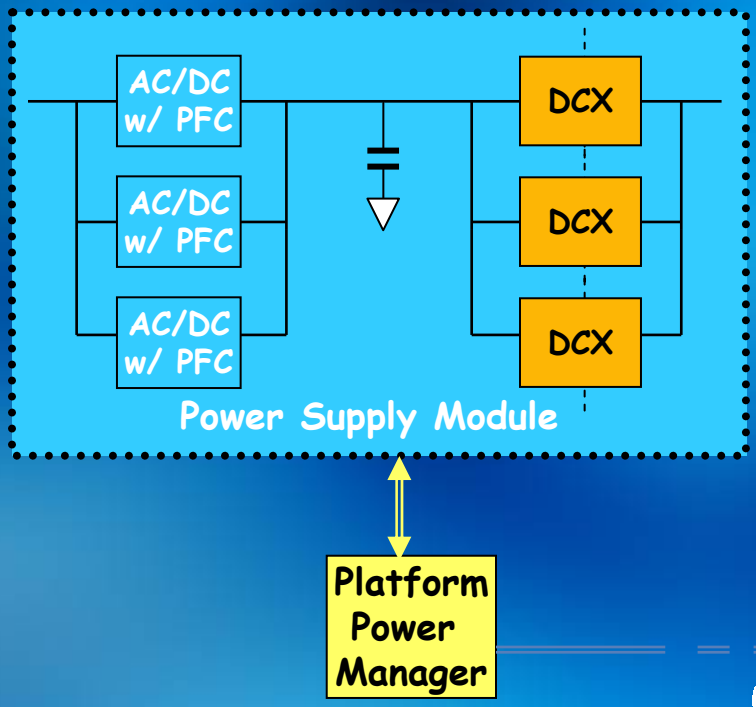

Load Adaptive PSU

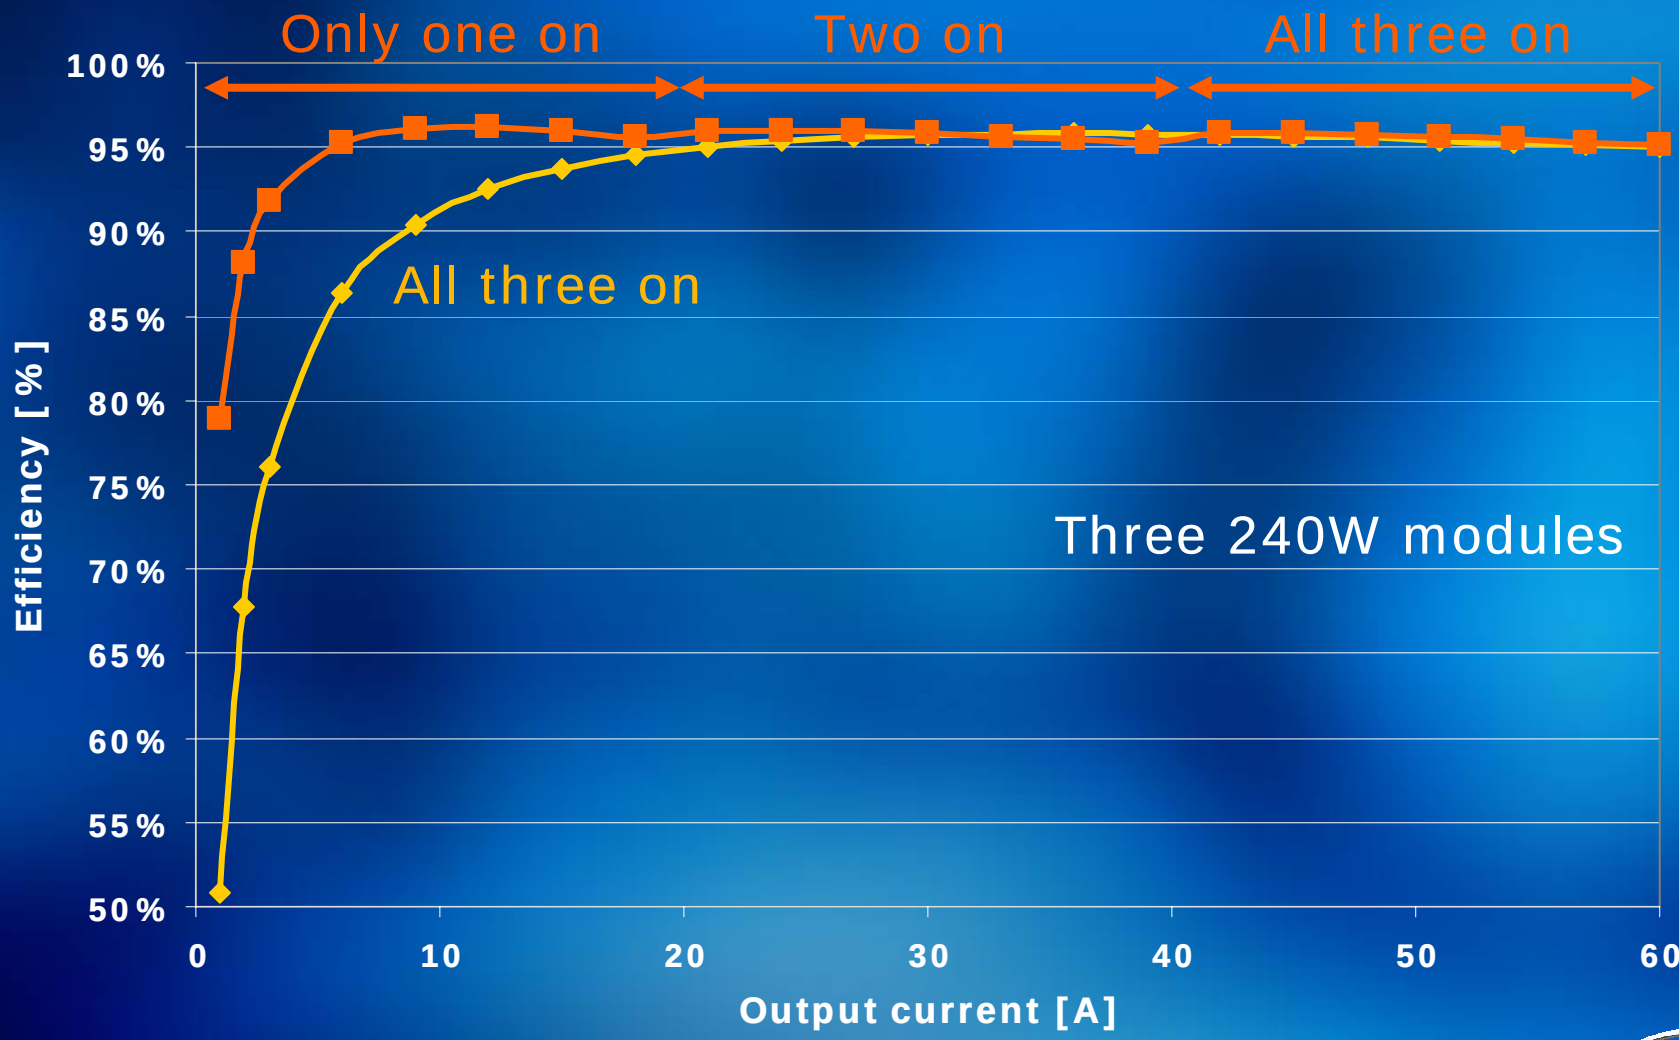

Demand Based Switching

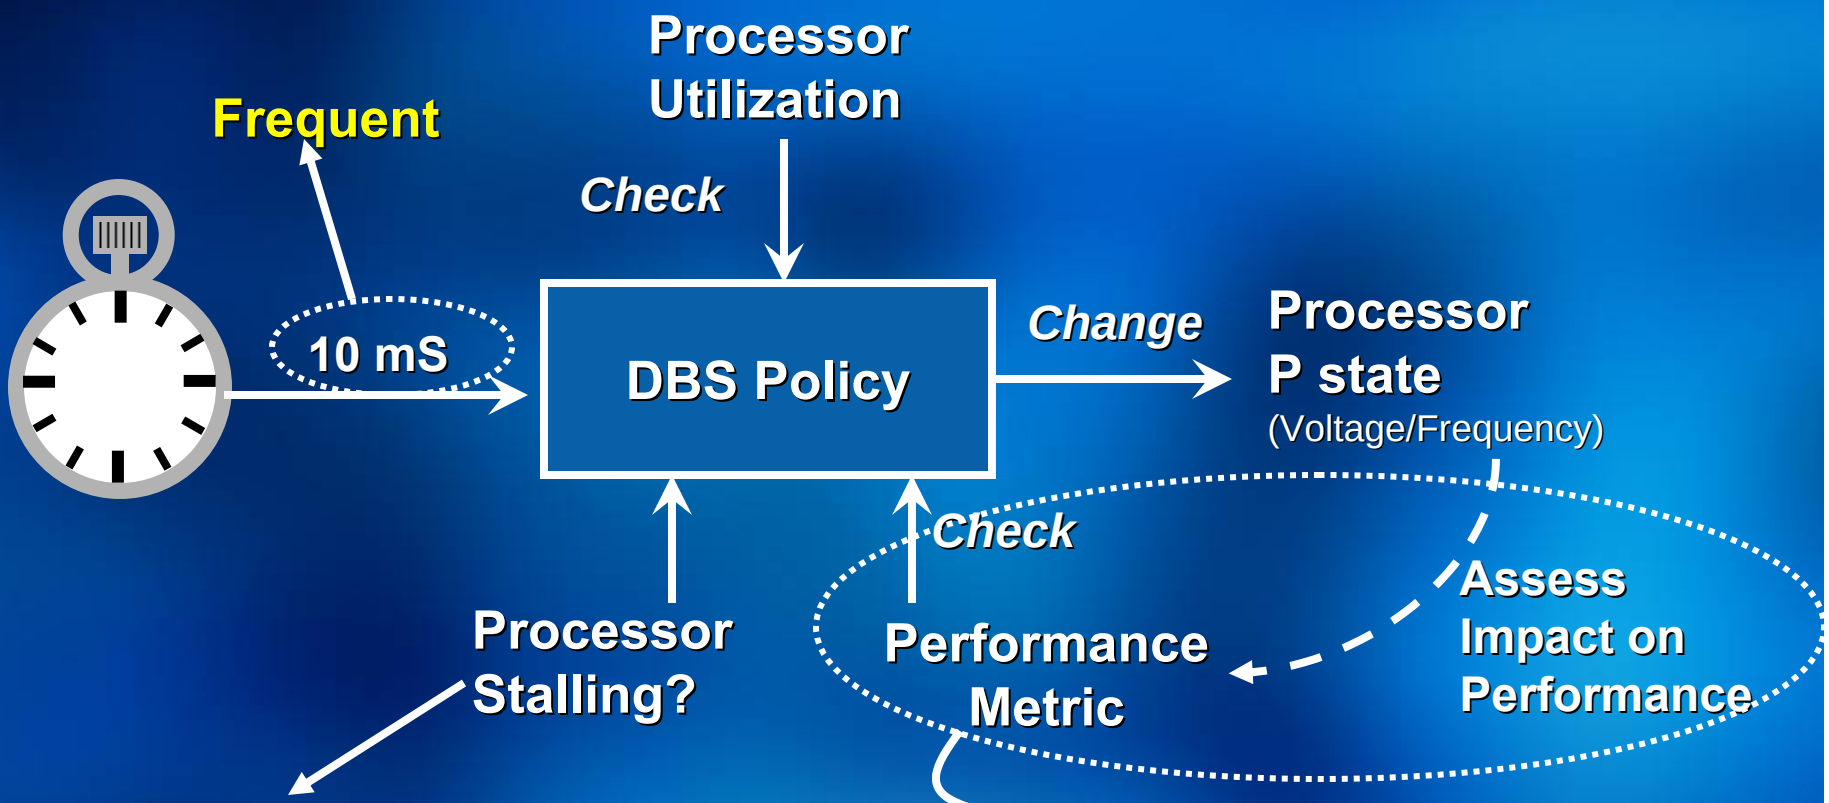

Reduce power often & also when processor stalls

Fine-grain Power Management

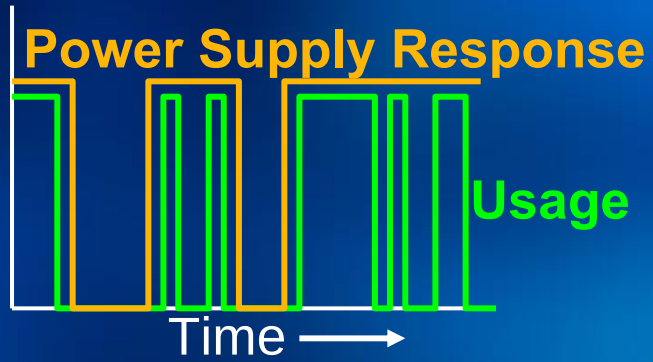

Improve response time
Bring power supply closer

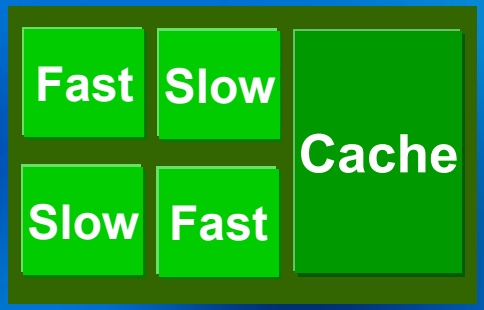

Provide multiple supply voltages
Fine grain Vdd and Freq scaling

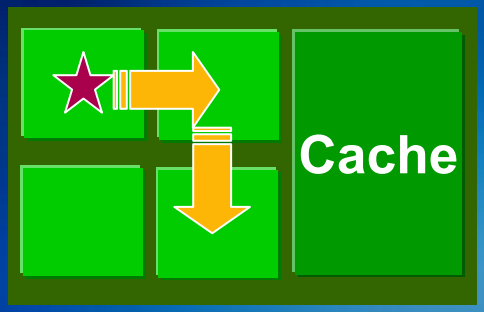

Core Hopping for hot spot & power density management

Summary

- Energy Efficiency will continue to be a major design goal
- Holistic approach needed
  - Semiconductor process technoloby
  - Circuit design
  - Micro-architecture
  - Platform architecture
  - Power supply design
  - Software


intel®
Leap ahead™